SYNTHESIZING NESTED RELATIONAL QUERIES FROM IMPLICIT SPECIFICATIONS: VIA MODEL THEORY AND VIA PROOF THEORY

MICHAEL BENEDIKT [a], CÉCILIA PRADIC [b], AND CHRISTOPH WERNHARD [c]

[a] University of Oxford, UK
   e-mail address: michael.benedikt@cs.ox.ac.uk

[b] University of Swansea, UK
   e-mail address: c.pradic@swansea.ac.uk

[c] University of Potsdam, Germany
   e-mail address: christoph.wernhard@uni-potsdam.de

ABSTRACT. Derived datasets can be defined implicitly or explicitly. An *implicit definition* (of dataset O in terms of datasets $\vec{I}$) is a logical specification involving two distinguished sets of relational symbols. One set of relations is for the "source data" $\vec{I}$, and the other is for the "interface data" O. Such a specification is a valid definition of O in terms of $\vec{I}$, if any two models of the specification agreeing on $\vec{I}$ agree on O. In contrast, an *explicit definition* is a transformation (or "query" below) that produces O from $\vec{I}$. Variants of *Beth's theorem* [Bet53] state that one can convert implicit definitions to explicit ones. Further, this conversion can be done effectively given a proof witnessing implicit definability in a suitable proof system.

We prove the analogous implicit-to-explicit result for nested relations: implicit definitions, given in the natural logic for nested relations, can be converted to explicit definitions in the nested relational calculus (NRC). We first provide a model-theoretic argument for this result, which makes some additional connections that may be of independent interest, between NRC queries, *interpretations*, a standard mechanism for defining structure-to-structure translation in logic, and between interpretations and implicit to definability "up to unique isomorphism". The latter connection uses a variation of a result of Gaifman concerning "relatively categorical" theories. We also provide a proof-theoretic result that provides an effective argument: from a proof witnessing implicit definability, we can efficiently produce an NRC definition. This will involve introducing the appropriate proof system for reasoning with nested sets, along with some auxiliary Beth-type results for this system. As a consequence, we can effectively extract rewritings of NRC queries in terms of NRC views, given a proof witnessing that the query is determined by the views.

1. Introduction

One way of describing a virtual datasource is via *implicit definition*: a specification Σ – e.g. in logic – involving symbols for the "virtual" object O and the stored "input" data $\vec{I}$. The specification may mention other data objects (e.g. auxiliary views). But to be an implicit definition, any two models of Σ that agree on $\vec{I}$ must agree on O. In the case where Σ is in first-order logic, this hypothesis can be expressed as a first-order entailment, using two copies of the vocabulary, primed and unprimed, representing the two models:

$$\Sigma \wedge \Sigma' \wedge \bigwedge_{I_i \in \vec{I}} \forall \vec{x}_i \; [I_i(\vec{x}_i) \leftrightarrow I'_i(\vec{x}_i)] \models \forall \vec{x} \; [O(\vec{x}) \leftrightarrow O'(\vec{x})] \qquad (\star)$$

Above Σ' is a copy of Σ with primed versions of each predicate.

A fundamental result in logic states that we can replace an implicit definition with an *explicit definition*: a first-order query Q such that whenever $\Sigma(\vec{I}, O, \ldots)$ holds, $O = Q(\vec{I})$. The original result of this kind is *Beth's theorem* [Bet53], which deals with classical first-order logic. Segoufin and Vianu's [SV05] looks at the case where Σ is in *active-domain first-order logic*, or equivalently a Boolean relational algebra expression. Their conclusion is that one can produce an explicit definition of O over $\vec{I}$ in relational algebra. [SV05] focused on the special case where $\Sigma(I_1 \ldots I_j, \vec{B}, O)$ specifies each I_i as a view defined by an active-domain first-order formula φ_{V_i} over base data $\vec{B}$, and also defines O as an active-domain first-order query φ_Q over $\vec{B}$. In this case, Σ implicitly defining O in terms of $\vec{I}$ is called "determinacy of the query by the views". Segoufin and Vianu's result implies that *whenever a relational algebra query Q is determined by relational algebra views $\vec{V}$, then Q is rewritable over the views by a relational algebra query.*

Prior Beth-style results like [Bet53, SV05] are effective. From a proof of the entailment $(\star)$ in a suitable proof system, one can extract an explicit definition effectively, even in polynomial time. While in early proofs of Beth's theorem, the proof systems were custom-designed for the task of proving implicit definitions, and the bounds were not stated, later on standard proof systems such as tableaux [Smu68b] or resolution [Hua95] were employed, and the polynomial claim was explicit. It is important that in our definition of implicit definability, we require the existence of a proof witness. By the completeness theorem for first-order logic, requiring such a proof witness is equivalent to demanding that implicit definability of O over $\vec{I}$ holds for all instances, not just finite ones.

Nested relations are a natural data model for hierarchical data. Nested relations are objects within a type system built up from basic types via tupling and a set-former. In the 1980's and 90's, a number of algebraic languages were proposed for defining transformations on nested collections. Eventually a standard language emerged, the *nested relational calculus* (NRC). The language is strongly-typed and functional, with transformations built up via tuple manipulation operations as well as operators for lifting transformations over a type T to transformations taking as input a set of objects of type T, such as singleton constructors and a mapping operator. One common formulation of these uses variables and a "comprehension" operator for forming new objects from old ones [BNTW95], while an alternative algebraic formalism presents the language as a set of operators that can be freely composed. It was shown that each NRC expression can be evaluated in polynomial time in the size of a finite data input, and that when the input and output is "flat" (i.e. only one level of nesting), NRC expresses exactly the transformations in the standard relational database

language relational algebra. Wong's thesis [Won94] summarizes the argument made by this line of work "NRC can be profitably regarded as the 'right' core for nested relational languages". NRC has been the basis for most work on transforming nested relations. It is the basis for a number of commercial tools [MGL$^+$10], including those embedding nested data transformations in programming languages [MBB06], in addition to having influence in the effective implementation of data transformations in functional programming languages [GHHW18, Gib16].

Although NRC can be applied to other collection types, such as bags and lists, we will focus here on just nested sets. We will show a new connection between NRC and first-order logic.

There is a natural logic for describing properties of nested relations, the well-known Δ_0 *formulas*, built up from equalities using quantifications $\exists x \in \tau$ and $\forall y \in \tau$ where τ is a term. For example, formula $\forall x \in c\ \pi_1(x) \in \pi_2(x)$ might describe a property of a nested relation c that is a set of pairs, where the first component of a pair is of some type T and the second component is a set containing elements of type T. A Δ_0 formula $\Sigma(\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k, \mathsf{o}_{out})$ over variables $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$ and variable o_{out} thus defines a relationship between $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$ and o_{out}. For such a formula to define a transformation it must be *functional*: it must enforce that o_{out} is determined by the values of $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$. More generally, if we have a formula $\Sigma(\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k, \mathsf{o}_{out}, \vec{a})$, we say that Σ *implicitly defines* o_{out} *as a function of* $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$ if:

> *For each two bindings σ_1 and σ_2 of the variables $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k, \vec{a}, \mathsf{o}_{out}$ to nested relations satisfying Σ, if σ_1 and σ_2 agree on each o_{in}^i, then they agree on o_{out}.* $\qquad(1.1)$

That is, Σ entails that the value of o_{out} is a partial function of the value of $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$.

Note that when we say "for each binding of variables to nested relations" in the definitions above, we include infinite nested relations as well as finite ones. An alternative characterization of Σ being an implicit definition, which will be more relevant to us in the second part of the paper, is that there is a proof that Σ defines a functional relationship. Note that (1.1) could be expressed as a first-order entailment: $\Sigma(\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k, \mathsf{o}_{out}, \vec{a}) \wedge \Sigma(\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k, \mathsf{o}_{out}', \vec{a}') \models \mathsf{o}_{out} = \mathsf{o}_{out}'$ where in the entailment we omit some first-order "sanity axioms" about tuples and sets. We refer to a proof of (1.1) for a given Σ and subset of the input variables $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$, as a *proof that Σ implicitly defines o_{out} as a function of $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$*, or simply a *proof of functionality* dropping Σ, o_{out}, and $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$ when they are clear from context. By the completeness theorem of first-order logic, whenever Σ defines o_{out} as a function of $\mathsf{o}_{in}^1 \dots \mathsf{o}_{in}^k$ according to the semantic definition above, this could be witnessed by a proof in any of the standard complete proof calculi for classical first-order logic (e.g. tableaux, resolution). Such a proof will use the sanity axioms referred to above, which capture extensionality of sets, the compatibility of the membership relation with the type hierarchy, and properties of projections and tupling. This notion of proof is only presented as an illustration. In the second half of the paper, rather than using a general-purpose first-order system, we will present more restrictive proof calculi that are tailored to reasoning about equivalence of nested sets relative to Δ_0 theories.

Example 1.1. We consider a specification in logic involving two nested collections, F and G. The collection F is of type $\mathsf{Set}(\mathfrak{U} \times \mathfrak{U})$, where $\mathfrak{U}$ refers to the basic set of elements, the "Ur-elements" in the sequel. That is, F is a set of pairs. The collection G is of type $\mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$, a set whose members are pairs, the first component an element and the second a set.

Our specification Σ will state that for each element g in G there is an element f_1 appearing as the first component of a pair in F, such that g represents f_1, in the sense that its first component is f_1 and its second component accumulates all elements paired with f_1 in F. This can be specified easily by a Δ_0 formula:

$$\forall g \in G \ \exists f \in F \quad \pi_1(g) = \pi_1(f) \ \wedge \ \forall x \in \pi_2(g) \ \langle \pi_1(f), x \rangle \in F$$
$$\wedge \ \forall f' \in F \ \left[\pi_1(f') = \pi_1(f) \rightarrow \pi_2(f') \in \pi_2(g) \right]$$

Σ also states that for each element f_1 lying within a pair in F there is a corresponding element g of G that pairs f_1 with all of the elements linked with f in F.

$$\forall f \in F \ \exists g \in G \quad \pi_1(g) = \pi_1(f) \ \wedge \ \forall x \in \pi_2(g) \ \langle \pi_1(f), x \rangle \in F$$
$$\wedge \ \forall f' \in F \ \left[\pi_1(f') = \pi_1(f) \rightarrow \ \pi_2(f') \in \pi_2(g) \right]$$

We can argue that in a nested relation satisfying Σ, G is a function of F. Thus Σ implicitly defines a function from F to G.

We give the argument informally here. Fixing F, G and F, G' satisfying Σ, we will prove that if $g \in G$ then $g \in G'$. The proof begins by using the conjunct in the first item to obtain an $f \in F$. We can then use the second item on G' to obtain a $g' \in G'$. We now need to prove that $g' = g$. Since g and g' are pairs, it suffices to show that their two projections are the same. We can easily see that $\pi_1(g) = \pi_1(f) = \pi_1(g')$, so it suffices to prove $\pi_2(g') = \pi_2(g)$. Here we will make use of extensionality, arguing for containments between $\pi_2(g')$ and $\pi_2(g)$ in both directions. In one direction we consider an $x \in \pi_2(g')$, and we need to show x is in $\pi_2(g)$. By the second conjunct in the second item we have $\langle \pi_1(f), x \rangle \in F$. Now using the first item we can argue that $x \in \pi_2(g)$. In the other direction we consider $x \in \pi_2(g)$, we can apply the first item to claim $\langle \pi_1(f), x \rangle \in F$ and then employ the second item to derive $x \in \pi_2(g')$.

Now let us consider G as the input and F as the output. We cannot say that Σ describes F as a *total* function of G, since Σ enforces constraints on G: that the second component of a pair in G cannot be empty, and that any two pairs in G that agree on the first component must agree on the second. But we can prove from Σ that F is a partial function of G: fixing F, G and F', G satisfying Σ, we can prove that $F = F'$.

Our first contribution is to show that whenever a Δ_0 formula Σ implicitly defines a function $\mathcal{T}$, that function can be expressed in a slight variant of NRC. The result can be seen as an analog of the well-known Beth definability theorem for first-order logic [Bet53].

The argument that we employ in our first contribution will go through some connections that give further insight. We first note that NRC-expressible functions have the same expressiveness as *interpretations*, a standard way of defining structure-to-structure transformations using logic. We will need a special notion of interpretation appropriate for nested sets. We will then establish an equivalence between interpretations and transformations that are *implicitly definable up to unique isomorphism*. This equivalence will hold in a much more general setting of multi-sorted first-order logic. The argument will be model-theoretic and non-constructive, relying on a variation of *Gaifman's coordinatisation theorem* [Hod93]. Putting these two connections together will establish our result.

Our second contribution is an effective version. We will start by providing a proof system which we show is complete for entailments involving Δ_0 formulas over nested relations. An advantage of our system compared to a classical first-order proof system mentioned above, is that we do not require special axioms about sets, like extensionality. In particular, we

never need to reason about formulas that are not Δ_0. We give two variations of the system, one that is lower-level (and less succinct). Our effective variant is then:

> From a proof p that Σ implicitly defines o in terms of $\vec{i}$, we can obtain, in PTIME, an NRC expression E that explicitly defines o from $\vec{i}$, relative to Σ. (1.2)

The PTIME claim refers to the lower-level calculus. The proof of the effective Beth result uses completely distinct techniques from the model-theoretic argument. We synthesize NRC expressions directly, rather than going through interpretations. The key to our proof-theoretic analysis is an auxiliary result about equalities between Δ_0 expressions, the NRC Parameter Collection Theorem, Theorem 5.9. It is a kind of interpolation theorem for parameterized definability. Roughly speaking, it says that when the conjunction of two Δ_0 formulas proves that two variables are equal, then there must be an expression definable with parameters that are common to both formulas that sits between them. We also make use of the interplay between reasoning about equivalence of nested sets "up to extensionality", without extensionality axioms, and the equality of nested sets in the presence of extensionality axioms.

A special case of our results will be in the case of NRC views and queries.

Example 1.2. We consider a variation of Example 1.1, where our specification $\Sigma(Q, V, B)$ describes a view V, a query Q, as well as some constraints on the base data B. Our base data B is of type $\mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$, where $\mathfrak{U}$ refers to the basic set of elements, the "Ur-elements". That is, B is a set of pairs, where the first item is a data item and the second is a set of data items. View V is of type $\mathsf{Set}(\mathfrak{U} \times \mathfrak{U})$, a set of pairs, given by the query that is the usual "flattening" of B: in NRC this can be expressed as $\{\langle \pi_1(b), c \rangle \mid c \in \pi_2(b) \mid b \in B\}$. The view definition can be converted to a specification in our logic.

A query Q might ask for a selection of the pairs in B, those whose first component is contained in the second: $\{b \in B \mid \pi_1(b) \in \pi_2(b)\}$. The definition of Q can also be incorporated into our specification.

View V is not sufficient to answer Q in general. This is the case if we assume as part of Σ an integrity constraint stating that the first component of B is a key. We can argue that $\Sigma(Q, V, B)$ implicitly defines Q in terms of V.

From first our main contribution, it follows that there is an NRC rewriting of Q in terms of V. From our second contribution, it follows that there is a polynomial time function that takes as input a proof in a certain proof system formalizing the implicit definability of Q in terms of V, which produces as output an NRC rewriting of Q in terms of V.

Overall our results show a close connection between logical specifications of transformations on nested collections and the functional transformation language NRC, a result which is not anticipated by the prior theory.

Organization. After a discussion of related work in Section 2, we provide preliminaries in Section 3. Section 4 presents our first main result, the expressive equivalence of implicit and explicit definitions, proven using model-theoretic techniques. As mentioned above, this requires an excursion into the relationship between NRC and a notion of first-order interpretation. Section 5 presents the effective version, which relies on a proof system for reasoning with nested relations. We close with discussion in Section 6. Some proofs of a routine nature, as well as some auxiliary results, are deferred to the appendix.

2. Related work

In addition to the theorems of Beth and Segoufin-Vianu mentioned in the introduction, there are numerous works on effective Beth-style results for other logics. Some concern fragments of classical first-order logic, such as the guarded fragment [HMO99, BtCV16]; others deal with non-classical logics such as description logics [tCFS13]. The Segoufin-Vianu result is closely related to variations of Beth's theorem and Craig interpolation for relativized quantification, such as Otto's interpolation theorem [Ott00]. There are also effective interpolation and definability results for logics richer than or incomparable to first-order logic, such as fragments of fixpoint logics [DH00, BBV19]. There are even Beth-style results for full infinitary logic [LE65], but there one can not hope for effectivity. The connection between Beth-style results and view rewriting originates in [SV05, NSV10]. The idea of using effective Beth results to generate view rewritings from proofs appears in [FKN13], and is explored in more detail first in [TW11] and later in [BCLT16].

Our first main result relates to Beth theorems "up-to-isomorphism". Our implicit definability hypothesis is that two models that satisfy a specification and agree on the inputs must agree on the output nested relations, where "agree on the output" means up to extensional equivalence of sets, which is a special (definable) kind of isomorphism. Beth-like theorems up to isomorphism originate in Gaifman's [Gai74]. are studied extensively by Hodges and his collaborators (e.g. [Hod75, HHM90, Hod93]). The focus of these works is model-theoretic, with emphasis on connections with categoricity and classification in classical model theory. More specifically, [Hod93] defines the notion of *rigidly relatively categorical* which is the single-sorted analog of the notion of implicitly interpretable which we will introduce in the model-theoretic part of this work. [Hod93] does not prove any connection of this notion to explicit interpretability, although he proves the equivalence with a related notion called "coordinatisability". Most of the ingredients in our main model-theoretic arguments are present in his exposition. The later unpublished draft [AMN08] extends these ideas to a multi-sorted setting, but without full proofs.

Conference versions. Our first contribution comes from the conference paper [BP21]. Our second result derives from the conference paper [BPW23].

3. Preliminaries

3.1. **Nested relations.** We deal with schemas that describe objects of various *types* given by the following grammar.

$$T,\ U ::= \mathfrak{U} \mid T \times U \mid \mathsf{Unit} \mid \mathsf{Set}(T)$$

For simplicity throughout the remainder we will assume only two basic types. There is the one-element type Unit, which will be used to construct Booleans. And there is $\mathfrak{U}$, the "scalars" or *Ur-elements* whose inhabitants are not specified further. From the Ur-elements and a unit type we can build up the set of types via product and the power set operation. We use standard conventions for abbreviating types, with the n-ary product abbreviating an iteration of binary products. A *nested relational schema* consists of declarations of variable names associated to objects of given types.

Example 3.1. An example nested relational schema declares two objects $R : \mathsf{Set}(\mathfrak{U} \times \mathfrak{U})$ and $S : \mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$. That is, R is a set of pairs of Ur-elements: a standard "flat" binary relation. S is a collection of pairs whose first elements are Ur-elements and whose second elements are sets of Ur-elements.

The types have a natural interpretation. The unit type has a unique member and the members of $\mathsf{Set}(T)$ are the sets of members of T. An *instance* of such a schema is defined in the obvious way.

For the schema in Example 3.1 above, assuming that $\mathfrak{U} = \mathbb{N}$, one possible instance has $R = \{\langle 4, 6 \rangle, \langle 7, 3 \rangle\}$ and $S = \{\langle 4, \{6, 9\} \rangle\}$.

3.2. Δ_0 formulas. We need a logic appropriate for talking about nested relations. A natural and well-known subset of first-order logic formulas with a set membership relation are the Δ_0 formulas. They are built up from equality of Ur-elements via Boolean operators as well as relativized existential and universal quantification. All terms involving tupling and projections are allowed. Our definition of Δ_0 formula is a variation of a well-studied notion in set theory [Jec03].

Formally, we deal with multi-sorted first-order logic, with sorts corresponding to each of our types. We use the following syntax for Δ_0 formulas and terms. Terms are built from variables using tupling and projections. All formulas and terms are assumed to be well-typed in the obvious way, with the expected sort of t and u being $\mathfrak{U}$ in expressions $t =_{\mathfrak{U}} u$ and $t \neq_{\mathfrak{U}} u$, and in $\exists t \in_T u \; \varphi$ the sort of t is T and the sort of u is $\mathsf{Set}(T)$.

$$
\begin{aligned}
t, u \quad &::= \quad x \mid \langle \rangle \mid \langle t, u \rangle \mid \pi_1(t) \mid \pi_2(t) \\
\varphi, \psi \quad &::= \quad t =_{\mathfrak{U}} t' \mid t \neq_{\mathfrak{U}} t' \mid \top \mid \bot \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \\
&\qquad \forall x \in_T t \; \varphi(x) \mid \exists x \in_T t \; \varphi(x)
\end{aligned}
$$

We call a formula *atomic* if it does not have a strict subformula. In the syntax presented above, atomic formulas are either of the form $t =_{\mathfrak{U}} t'$ or $t \neq_{\mathfrak{U}} t'$ (note that there are no equalities for sorts other than $\mathfrak{U}$). Negation $\neg\varphi$ will be defined as a macro by induction on φ by dualizing every connective and then the atomic formulas recursively. Other connectives can be derived in the usual way on top of negation: $\varphi \rightarrow \psi$ by $\neg\varphi \vee \psi$.

More crucial is the fact that membership is not itself a formula, it is only used in the relativized quantifiers. An *extended Δ_0 formula* allows also membership atomic formulas $x \in_T y$, $x \notin_T y$ and equalities $x =_T y$, $x \neq_T y$ at every type T.

The notion of an extended Δ_0 formula φ *entailing* another formula ψ is the standard one in first-order logic, meaning that every model of φ is a model of ψ. We emphasize here that by every model, we include models where membership is not extensional. We will require only one "sanity axiom": projection and tupling commute. An important point is that there is no distinction between entailment in such "general models" and entailment over nested relations.

Proposition 3.2. *For φ and ψ are Δ_0, rather than extended Δ_0, φ entails ψ iff every nested relation that models φ is a model of ψ.*

The point above is due to two facts. First, we have neither $\in$ or equality at higher types as predicates. This guarantees that any model can be modified, without changing the truth value of Δ_0 formulas, into a model satisfying extensionality: if we have x and y with $(\forall z \in_T x \; z \in_T y) \wedge (\forall z \in_T y \; z \in_T x)$ then x and y must be the same. Secondly, a well-typed extensional model is isomorphic to a nested relation, by the well-known Mostowski collapse

$$E, E' ::= \quad x \mid \langle\rangle \mid \langle E, E' \rangle \mid \pi_1(E) \mid \pi_2(E) \mid \qquad \text{(variable, (un)tupling)}$$
$$\{E\} \mid \text{GET}_T(E) \mid \bigcup\{E \mid x \in E'\} \quad \text{((un)nesting, binding union)}$$
$$\mid \emptyset \mid E \cup E' \mid E \setminus E' \qquad\qquad \text{(finite unions, difference)}$$

Figure 1: NRC syntax (typing rules omitted)

construction that iteratively identifies elements that have the same members. The lack of primitive membership and equality relations in Δ_0 formulas allows us to avoid having to consider extensionality axioms, which would require special handling in our proof system.

Equality, inclusion and membership predicates "up to extensionality" may be defined as macros by induction on the involved types, while staying within Δ_0 formulas. Formally we have:

Definition 3.3.

$$
\begin{aligned}
t \, \hat{\in}_T \, u &:= \exists z \in_T u \; t \equiv_T z \\
t \subseteq_T u &:= \forall z \in_T t \; z \, \hat{\in}_T \, u \\
t \equiv_{\mathsf{Set}(T)} u &:= t \subseteq_T u \; \wedge \; u \subseteq_T t \\
t \equiv_{\mathsf{Unit}} u &:= \top \\
t \equiv_{\mathfrak{A}} u &:= t =_{\mathfrak{A}} u \\
t \equiv_{T_1 \times T_2} u &:= \pi_1(t) \equiv_{T_1} \pi_1(u) \; \wedge \; \pi_2(t) \equiv_{T_2} \pi_2(u)
\end{aligned}
$$

We will use small letters for variables in Δ_0 formulas, except in examples when we sometimes use capitals to emphasize that an object is of set type. We drop the type subscripts T in bounded quantifiers, primitive memberships, and macros $\equiv_T$ when clear. Of course membership-up-to-equivalence $\hat{\in}$ and membership $\in$ agree on extensional models. But $\in$ and $\hat{\in}$ are not interchangeable on general models, and hence are not interchangeable in Δ_0 formulas. For example:

$$x \in y, x \in y' \models \exists z \in y \; z \in y'$$

But we do not have

$$x \, \hat{\in} \, y, x \, \hat{\in} \, y' \models \exists z \in y \; z \in y'$$

3.3. **Nested Relational Calculus.** We review the main language for declaratively transforming nested relations, Nested Relational Calculus (NRC). Variables occurring in expressions are typed, and each expression is associated with an *output type*, both of these being in the type system described above. We let Bool denote the type Set(Unit). Then Bool has exactly two elements, and will be used to simulate Booleans. The grammar for NRC expressions is presented in Figure 1.

The definition of the free and bound variables of an expression is standard, the union operator $\bigcup\{E \mid x \in R\}$ binding the variable x. The semantics of these expressions should be fairly evident, see [Won94]. If E has type T, and has input (i.e. free) variables $x_1 \ldots x_n$ of types $T_1 \ldots T_n$, respectively, then the semantics associates with E a function that given a binding associating each free variable with a value of the appropriate type, returns an object of type T. For example, the expression $\langle\rangle$ always returns the empty tuple, while $\emptyset_T$ returns the empty set of type T.

As explained in prior work (e.g. [Won94]), on top of the NRC syntax above we can support richer operations as "macros". For every type T there is an NRC expression $=_T$ of type Bool representing equality of elements of type T. In particular, there is an expression

$=_\mathfrak{U}$ representing equality between Ur-elements. For every type T there is an NRC expression $\in_T$ of type Bool representing membership between an element of type T in an element of type $\mathsf{Set}(T)$. We can define conditional expressions, joins, projections on k-tuples, and k-tuple formers. Arbitrary arity tupling and projection operations $\langle E_1, \ldots E_n \rangle$, $\pi_j(E)$ for $j > 2$ can be seen as abbreviations for a composition of binary operations. Further

- If B is an expression of type Bool and E_1, E_2 expressions of type T, then there is an expression $\mathsf{case}(B, E_1, E_2)$ of type T that implements "if B then E_1 else E_2".
- If E_1 and E_2 are expressions of type $\mathsf{Set}(T)$, then there are expressions $E_1 \cap E_2$ and $E_1 \setminus E_2$ of type $\mathsf{Set}(T)$.

The derivations of these are not difficult. For example, the conditional required by the first item is given by:

$$\bigcup \{\{E_1\} \mid x \in B\} \cup \bigcup \{\{E_2\} \mid x \in (\neg B)\}$$

Finally, we note that NRC is closed under Δ_0 *comprehension*: if E is in NRC, φ is an extended Δ_0 formula, then we can efficiently form an expression $\{z \in E \mid \varphi\}$ which returns the subset of E such that φ holds. We make use of these macros freely in our examples of NRC, such as Example 1.2.

Example 3.4. Consider an input schema including a binary relation $F : \mathsf{Set}(\mathfrak{U} \times \mathfrak{U})$. The query $\mathcal{T}_{\mathsf{Proj}}$ with input F returning the projection of F on the first component can be expressed in NRC as $\bigcup \{\{\pi_1(f)\} \mid f \in F\}$. The query $\mathcal{T}_{\mathsf{Filter}}$ with input F and also v of type $\mathfrak{U}$ that filters F down to those pairs which agree with v on the first component can be expressed in NRC as $\bigcup \{\mathsf{case}([\pi_1(f) =_\mathfrak{U} v], \{f\}, \emptyset) \mid f \in F\}$. Consider now the query $\mathcal{T}_{\mathsf{Group}}$ that groups F on the first component, returning an object of type $\mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$. The query can be expressed in NRC as $\bigcup \{\{\langle v, \bigcup \{\{\pi_2(f)\} \mid f \in \mathcal{T}_{\mathsf{Filter}}\}\rangle\} \mid v \in \mathcal{T}_{\mathsf{Proj}}\}$. Finally, consider the query $\mathcal{T}_{\mathsf{Flatten}}$ that flattens an input G of type $\mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$. This can be expressed in NRC as

$$\bigcup \left\{ \bigcup \{\{\langle \pi_1(g), x\rangle\} \mid x \in \pi_2(g)\} \mid g \in G \right\}$$

The language NRC as originally defined cannot express certain natural transformations whose output type is $\mathfrak{U}$. To get a canonical language for such transformations, above we included in our NRC syntax a family of operations $\mathrm{GET}_T : \mathsf{Set}(T) \to T$ that extracts the unique element from a singleton. GET was considered in [Won94]. The semantics are: if E returns a singleton set $\{x\}$, then $\mathrm{GET}_T(E)$ returns x; otherwise it returns some default object of the appropriate type. In [Suc95], it is shown that GET is not expressible in NRC at sort $\mathfrak{U}$. However, GET_T for general T is definable from $\mathrm{GET}_\mathfrak{U}$ and the other NRC constructs.

Since GET will be needed for our key results, in the remainder of the paper, *we will write simply* NRC *for* NRC *as defined as usual, augmented with* GET. The role of GET will only be for transformations that return something of type $\mathfrak{U}$.

3.4. Connections between NRC queries using Δ_0 formulas.

Since we have a Boolean type in NRC, one may ask about the expressiveness of NRC for defining transformations of shape $T_1, \ldots, T_n \to \mathsf{Bool}$. It turns out that they are equivalent to Δ_0 formulas. This gives one justification for focusing on Δ_0 formulas.

Proposition 3.5. *There is a polynomial time algorithm taking an extended Δ_0 formula $\varphi(\vec{x})$ as input and producing an NRC expression $\mathsf{Verify}_\varphi(\vec{x})$ of type Bool such that, for any valuation in any nested relation, $\mathsf{Verify}_\varphi(\vec{x})$ returns true if and only if $\varphi(\vec{x})$ holds.*

This useful result is proved by an easy induction over φ: see Appendix B for details.

In the opposite direction, given an NRC expression E with input relations $\vec{i}$, we can create a Δ_0 formula $\Sigma_E(\vec{i}, o)$ that is an *input-output specification of E*: a formula such that Σ_E implies $o = E(\vec{i})$ and whenever $o = E(\vec{i})$ holds there is a set of objects including $\vec{i}$ and o satisfying Σ_E. For the "composition-free" fragment – in which comprehensions $\bigcup$ can only be over input variables – this conversion can be done in PTIME. But it cannot be done efficiently for general NRC, under complexity-theoretic hypotheses [Koc06].

We also write entailments that use NRC expressions. For example, we write:

$$\varphi(x, \vec{c} \ldots) \models_{\mathsf{nested}} x \in E(\vec{c})$$

for φ Δ_0 and $E \in$ NRC. An entailment with $\models_{\mathsf{nested}}$ involving NRC expressions means that in every *nested relation* satisfying φ, x is in the output of E on $\vec{c}$. Note that the semantics of NRC expressions is only defined on nested relations.

4. Synthesizing via model theory: the expressive equivalence of NRC, interpretations, and implicit definitions

Our first result will show the expressive equivalence of several specification languages for transformations. We will show that NRC expressions are equivalent to implicit definitions, but in the process we will show that both transformation languages are equivalent to transformations given in a natural logical language which we call Δ_0 interpretations. While the equivalence between interpretations and NRC will be effective, the key direction from implicit definitions to interpretations will be a model-theoretic argument.

4.1. **Statement of the first result: equivalence between implicit and explicit.** We consider an input schema $\mathcal{SCH}_{in}$ with one input object o_{in} and an output schema with one output object o_{out}. Using product objects, we can easily model any nested relational transformation in this way. We deal with a Δ_0 formula $\varphi(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ with distinguished variables $\mathsf{o}_{in}, \mathsf{o}_{out}$. Recall from the introduction that such a formula *implicitly defines o_{out} as a function of o_{in}* if for each nested relation o_{in} there is at most one o_{out} such that $\varphi(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ holds for some $\vec{a}$. A formula $\varphi(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ *implicitly defines a function $\mathcal{T}$ from o_{in} to o_{out}* if for each o_{in}, $\varphi(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ holds for some $\vec{a}$ if and only $\mathcal{T}(\mathsf{o}_{in}) = \mathsf{o}_{out}$.

Example 4.1. Consider the transformation $\mathcal{T}_{\mathsf{Group}}$ from Example 3.4. It has a simple implicit Δ_0 definition, which we can restate as follows. First, define the auxiliary formula $\chi(x, p, R)$ stating that $\pi_1(p)$ is x and $\pi_2(p)$ is the set of the ys such that $\langle x, y \rangle$ is in R (the "fiber of R above x"):

$$\chi(x, p, R) := \pi_1(p) = x \wedge \left(\forall t' \in R \; [\pi_1(t') = x \; \rightarrow \; \pi_2(t') \in \pi_2(p)] \right) \wedge \forall z \in \pi_2(p) \; \langle x, z \rangle \in R$$

Then the transformation T_{Group} from Q to R is implicitly defined by

$$\forall t \in R \; \exists p \in Q \; \chi(\pi_1(t), p, R)) \wedge \forall p \in Q \; \chi(\pi_1(p), p, R)$$

We can now state our first main result:

Theorem 4.2 (Implicit-to-explicit for nested relations via model theory)**.** *For any Δ_0 formula $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ which implicitly defines o_{out} as a function of o_{in}, there is an NRC expression E such that whenever $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ holds, then $E(\mathsf{o}_{in}) = \mathsf{o}_{out}$.*

In particular, suppose that, in addition, for each o_{in} there is some o_{out} and $\vec{a}$ such that $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ holds: then the expression and the formula define the same transformation.

Recall that our notion of implicit definitions allows extra parameters $\vec{a}$. Sometimes these are called "projective" implicit definitions in the literature. From Theorem 4.2 we easily see that no additional expressiveness is gained by allowing parameters:

Corollary 4.3. *The following are equivalent for a transformation $\mathcal{T}$:*
(1) *$\mathcal{T}$ is implicitly definable by a Δ_0 formula*
(2) *$\mathcal{T}$ is implicitly definable by a Δ_0 formula $\varphi(\mathsf{o}_{in}, \mathsf{o}_{out})$*
(3) *$\mathcal{T}$ is NRC definable*

Finite instances versus all instances. In Theorem 4.2 and Corollary 4.3 we emphasize that our results concern the class $\mathsf{Fun}_{\mathsf{All}}$ of transformations $\mathcal{T}$ such that there is a Δ_0 formula Σ which defines a functional relationship between o_{in} and o_{out} on all instances, finite and infinite, and where the function agrees with $\mathcal{T}$. We can consider $\mathsf{Fun}_{\mathsf{All}}$ as a class of transformations on all instances or on all finite instances, but the class is defined by reference to all instances for o_{in}. Expressed semantically

$$\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a}) \wedge \Sigma(\mathsf{o}_{in}, \mathsf{o}'_{out}, \vec{a}') \models \mathsf{o}'_{out} = \mathsf{o}_{out}$$

An equivalent characterization of $\mathsf{Fun}_{\mathsf{All}}$ is *proof-theoretic*: these are the transformations such that there is a classical proof of functionality in a complete first-order proof system. There are various choices for the system. The most straightforward choice would be a system using some basic axioms about Ur-elements, products and projection functions, and the extensionality axiom for the membership relation. We will see a slightly different approach to proof systems in Section 5.

Whether one thinks of $\mathsf{Fun}_{\mathsf{All}}$ semantically or proof-theoretically, our results say that $\mathsf{Fun}_{\mathsf{All}}$ is identical with the set of transformations given by NRC expressions. But the proof-theoretic perspective is crucial in order to even talk about an effective synthesis procedure.

It is natural to ask about the analogous class $\mathsf{Fun}_{\mathsf{Fin}}$ of transformations $\mathcal{T}$ over *finite inputs* for which there is a Δ_0 $\Sigma_{\mathcal{T}}$ which is functional, when only finite inputs are considered, and where the corresponding function agrees with $\mathcal{T}$. It is well-known that $\mathsf{Fun}_{\mathsf{Fin}}$ is not identical to NRC and is not so well-behaved. The transformation returning the powerset of a given input relation o_{in} is in $\mathsf{Fun}_{\mathsf{Fin}}$: the powerset of a finite input o_{in} is the unique collection o_{out} of subsets of o_{in} that contains the empty set and such that for each element e of o_{in}, if a set s is in o_{out} then $s - \{e\}$ and $s \cup \{e\}$ are in o_{out}. From this we can see that $\mathsf{Fun}_{\mathsf{Fin}}$ contains transformations of high complexity. Indeed, even when considering transformations from flat relations to flat relations, $\mathsf{Fun}_{\mathsf{Fin}}$ contains transformations whose membership in polynomial time would imply that $\mathsf{UP} \cap \mathsf{coUP}$, the class of problems such that both the problem and its complement can be solved by an unambiguous non-deterministic polynomial time machine, is identical to PTIME [Kol90]. Most importantly for our goals, membership in $\mathsf{Fun}_{\mathsf{Fin}}$ is *not* witnessed by proofs in any effective proof system, since this set is not computably enumerable.

Total versus partial functions. When we have a proof that $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ defines o_{out} as a function of o_{in}, the corresponding function may still be partial. Our procedure will synthesize an expression E defining a total function that agrees with the partial function defined by Σ. If $\vec{a}$ is empty, we can also synthesize a Boolean NRC expression $\mathsf{Verify}_{\mathsf{InDomain}}$

that verifies whether a given o_{in} is in the domain of the function: that is whether there is o_{out} such that $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out})$ holds. $\mathsf{Verify}_{\mathsf{InDomain}}$ can be taken as:

$$\bigcup \{\mathsf{Verify}_{\Sigma}(\mathsf{o}_{in}, e) \mid e \in \{E(\mathsf{o}_{in})\}\}$$

where Verify_{Σ} is from Proposition 3.5.

When $\vec{a}$ is not empty we can not generate a domain check $\mathsf{Verify}_{\mathsf{InDomain}}$, since the auxiliary parameters might enforce some second-order property of $\vec{i}$: for example $\Sigma(i_0, i_1, a, o)$ might state that a is a bijection from i_0 to i_1 and $o = \langle i_1, i_2 \rangle$. This clearly defines a functional relationship between i_1, i_2 and o, but the domain consists of i_1, i_2 that have the same cardinality, which cannot be expressed in first-order logic.

Organization of the proof of the theorem. Our proof of Theorem 4.2 will go through a notion of Δ_0 interpretation, which we introduce in Subsection 4.2. We will show that Δ_0 interpretations define the same transformations as NRC, which will allow us to restate the main result as moving from implicit definitions to interpretations. We then proceed first by some reductions (Subsection 4.3), showing that it suffices to prove a general result about implicit definability and definability by interpretations in multi-sorted first-order logic, rather than dealing with higher-order logic and Δ_0 formulas. In Subsection 4.4 we give the argument for this multi-sorted logic theorem.

Interpolation for Δ_0 formulas. Often a key ingredient in moving from implicit to explicit definition is an *interpolation theorem*, stating that for each entailment between formulas φ_L and φ_R there is an intermediate formula (an *interpolant* for the entailment), which is entailed by φ_L and entails φ_R while using only symbols common to φ_L and φ_R. We can show using any of the standard approaches to interpolation (e.g. [Fit96]) that Δ_0 formulas admit interpolation.

Proposition 4.4. *Let Γ_L, Γ_R, and ψ be Δ_0 formulas and call $C = \mathsf{FV}(\Gamma_L) \cap (\mathsf{FV}(\Gamma_R) \cup \mathsf{FV}(\psi))$ the set of common free variables of Γ_L on the one hand and Γ_R or ψ on the other hand. If we have an entailment*

$$\Gamma_L, \; \Gamma_R \models \psi$$

then there exists a Δ_0 formula θ with $\mathsf{FV}(\theta) \subseteq C$ such that the following holds

$$\Gamma_L \models \theta \qquad and \qquad \Gamma_R, \theta \models \psi$$

A stronger effective statement – stating that the interpolant can be found efficiently in the size of a proof of the entailment – will be proven later in the paper: see Theorem 5.4.

The interpolation result above should be thought of as giving us the result we want for implicit definitions of *Boolean* variables. From it we can derive that whenever we have a Δ_0 $\Sigma(\vec{i} \ldots o)$ which implicitly defines a Boolean variable o in terms of input variables $\vec{i}$, there must be a Δ_0 $\varphi(\vec{i}, o)$ that defines o in terms of $\vec{i}$. Setting o to be true we get a formula $\varphi'(\vec{i})$ that defines the inputs that correspond to true. By Proposition 3.5, there is an NRC expression outputting a Boolean that explicitly defines o.

4.2. **Interpretations and nested relations.** Our first goal will be to show that for any Δ_0 implicit definitions there is an NRC query that realizes it. For this result, it will be useful to have another characterization of NRC, an equivalence with transformations defined by *interpretations.*

We first review the notion of an interpretation, which has become a common way of defining transformations using logical expressions [BDK18, CL07]. Let $\mathcal{SCH}_{in}$ and $\mathcal{SCH}_{out}$ be multi-sorted vocabularies. A first-order interpretation with input signature $\mathcal{SCH}_{in}$ and output signature $\mathcal{SCH}_{out}$ consists of:

- for each output sort S', a sequence of input sorts $\tau(\mathsf{S}') = \vec{\mathsf{S}}$,
- a formula $\varphi_{\underline{\equiv}}^{\mathsf{S}'}(\vec{x}_1, \vec{x}_2)$ for each output sort S' in $\mathcal{SCH}_{out}$ (where both tuples of variables $\vec{x}_1$ and $\vec{x}_2$ have types $\tau(\mathsf{S}')$),
- a formula $\varphi_{\mathsf{Domain}}^{\mathsf{S}'}(\vec{x}_1)$ for each output sort S' in $\mathcal{SCH}_{out}$ (where the variables $\vec{x}_1$ have types $\tau(\mathsf{S}')$),
- a formula $\varphi_R(\vec{x}_1, \ldots \vec{x}_n)$ for every relation R of arity n in $\mathcal{SCH}_{out}$ (where the variables $\vec{x}_i$ have types $\tau(\mathsf{S}'_i)$, provided the i^{th} argument of R has sort S'_i),
- for every function symbol $f(x_1, \ldots, x_k)$ of $\mathcal{SCH}_{out}$ with output sort S' and input x_i of sort S_i, a sequence of terms $\overline{f}_1(\vec{x}_1, \ldots, \vec{x}_k), \ldots, \overline{f}_m(\vec{x}_1, \ldots, \vec{x}_k)$ with sorts enumerating $\tau(\mathsf{S}_{out})$ (so in particular m corresponds to the length of $\tau(\mathsf{S}_{out})$) and $\vec{x}_i$ of sorts $\tau(\mathsf{S}_i)$.

subject to the following constraints

- $\varphi_{\underline{\equiv}}^{\mathsf{S}}(\vec{x}, \vec{y})$ should define a partial equivalence relation, i.e. be symmetric and transitive,
- $\varphi_{\mathsf{Domain}}^{\mathsf{S}}(\vec{x})$ should be equivalent to $\varphi_{\underline{\equiv}}^{\mathsf{S}}(\vec{x}, \vec{x})$,
- $\varphi_R(\vec{x}_1, \ldots, \vec{x}_n)$ and $\varphi_{\underline{\equiv}}^{\mathsf{S}_i}(\vec{x}_i, \vec{y}_i)$ for $1 \leq 1 \leq n$ (where S_i is the output sort associated with position i of the relation R) should jointly imply $\varphi_R(\vec{y}_1, \ldots, \vec{y}_n)$.
- for every output function symbol $f(x_1, \ldots, x_k)$ represented by terms $\vec{\overline{f}}(\vec{x})$, we have

$$\forall \vec{\vec{x}}\, \vec{\vec{y}}\, \left(\bigwedge_i \varphi_{\underline{\equiv}}^{\mathsf{S}_i}(\vec{x}_i, \vec{y}_i) \rightarrow \varphi_{\underline{\equiv}}^{\mathsf{S}'}(\vec{\overline{f}}(\vec{\vec{x}}), \vec{\overline{f}}(\vec{\vec{y}})) \right)$$

where S' is the sort of the output of f and the S_i correspond to the arities.

In $\varphi_{\underline{\equiv}}^{\mathsf{S}}$ and $\varphi_{\mathsf{Domain}}^{\mathsf{S}}$, each $\vec{x}_1, \vec{x}_2$ is a tuple containing variables of sorts agreeing with the prescribed sequence of input sorts for S'. Given a structure M for the input sorts and a sort S we call a binding of these variables to input elements of the appropriate input sorts an M, S *input match*. If in output relation R position i is of sort S_i, then in $\varphi_R(\vec{t}_1, \ldots \vec{t}_n)$ we require $\vec{t}_i$ to be a tuple of variables of sorts agreeing with the prescribed sequence of input sorts for S_i. Each of the above formulas is over the vocabulary of $\mathcal{SCH}_{in}$.

An interpretation $\mathcal{I}$ defines a function from structures over vocabulary $\mathcal{SCH}_{in}$ to structures over vocabulary $\mathcal{SCH}_{out}$ as follows:

- The domain of sort S' is the set of equivalence classes of the partial equivalence relation defined by $\varphi_{\underline{\equiv}}^{\mathsf{S}'}$ over the M, S' input matches.
- A relation R in the output schema is interpreted by the set of those tuples $\vec{a}$ such that $\varphi_R(\vec{t}_1, \ldots \vec{t}_n)$ holds for some $\vec{t}_1 \ldots \vec{t}_n$ with each $\vec{t}_i$ a representative of a_i.

An interpretation $\mathcal{I}$ also defines a map $\varphi \mapsto \varphi^*$ from formulas over $\mathcal{SCH}_{out}$ to formulas over $\mathcal{SCH}_{in}$ in the obvious way. This map commutes with all logical connectives and thus preserves logical consequence.

In the sequel, we are concerned with interpretations preserving certain theories consisting of sentences in first-order logic. Recall that a *theory* in first-order logic is a deductively closed

set of sentences. A sentence belonging to a given theory is called one of its *theorems*. Given a theory Σ over $\mathcal{SCH}_{in}$ and a theory Σ' over $\mathcal{SCH}_{out}$, we say that $\mathcal{I}$ is an interpretation of Σ' within Σ if $\mathcal{I}$ is an interpretation such that for every theorem φ of Σ', φ^* is a theorem of Σ. Since $\varphi \mapsto \varphi^*$ preserves logical consequence, if Σ' is generated by a set of axioms A, it suffices to check that Σ' proves φ for $\varphi \in A$.

Finally, we are also interested in interpretations restricting to the identity on part of the input. Suppose that $\mathcal{SCH}_{out}$ and $\mathcal{SCH}_{in}$ share a sort S. An interpretation $\mathcal{I}$ of $\mathcal{SCH}_{out}$ within $\mathcal{SCH}_{in}$ is said to preserve S if the output sort associated to S is S itself and the induced map of structures is the identity over S. Up to equivalence, that means we fix $\varphi_{\mathsf{Domain}}^T(x)$ to be, up to equivalence, $\top$, $\varphi_{\underset{=}{\mathsf{S}}}(x,y)$ to be the equality $x = y$ and map constants of type S to themselves.

Interpretations defining nested relational transformations. We now consider how to define nested relational transformations via interpretations. The main idea will be to restrict all the constituent formulas to be Δ_0 and to relativize the notion of interpretation to a background theory that corresponds to our sanity axioms about tupling and sets.

We define the notion of *subtype* of a type T inductively as follows.

- T is a subtype of $\mathsf{Set}(T')$ if $T = \mathsf{Set}(T')$ or if it is a subtype of T'.
- T is a subtype of $T_1 \times T_2$ if $T = T_1 \times T_2$ or if it is a subtype of either T_1 or T_2.
- The only subtypes of $\mathfrak{U}$ and Unit are themselves.

For every type T, we build a multi-sorted vocabulary $\mathcal{SCH}_T$ as follows.

- The sorts are all subtypes of T, Unit and $\mathsf{Bool} = \mathsf{Set}(\mathsf{Unit})$.
- The function symbols are the projections, tupling, the unique element of type Unit, the constants ff, tt of sort Bool representing $\emptyset, \{\langle\rangle\}$ and a special constant o of sort T.
- The relation symbols are the equalities at every sort and the membership predicates $\in_T$.

Let T_{obj} be a type which will represent the type of a complex object obj. We build a theory $\Sigma(T_{\mathsf{obj}})$ on top of $\mathcal{SCH}_{T_{\mathsf{obj}}}$ from the following axioms:

- Equality should satisfy the congruence axioms for every formula φ

$$\forall xy \ (x = y \ \wedge \ \varphi \ \rightarrow \ \varphi[y/x])$$

Note that it is sufficient to require this for atomic formulas to infer it for all formulas.
- We require that projection and tupling obey the usual laws for every type of $\mathcal{SCH}_{T_{\mathsf{obj}}}$.

$$\forall x^{T_1} \ y^{T_2} \ \pi_1(\langle x,y \rangle) = x \qquad \forall x^{T_1} \ y^{T_2} \ \pi_2(\langle x,y \rangle) = y \qquad \forall x^{T_1 \times T_2} \ \langle \pi_1(x), \pi_2(x) \rangle = x$$

- We require that Unit be a singleton and every $\mathsf{Set}(T)$ in $\mathcal{SCH}_{T_{\mathsf{obj}}}$

$$\forall x^{\mathsf{Unit}} \ \langle\rangle = x$$

- Lastly our theory imposes set extensionality

$$\forall x^{\mathsf{Set}(T)} \ y^{\mathsf{Set}(T)} \ \left([\forall z^T \ (z \in_T x \leftrightarrow z \in_T y)] \rightarrow x =_T y \right)$$

Note that in interpretations we associate the input to a structure that includes a distinguished constant. For example, an input of type $\mathsf{Set}(\mathfrak{U})$ will be coded by a structure with an element relation, an Ur-element sort, and a constant whose sort is the type $\mathsf{Set}(\mathfrak{U})$. In other contexts, like NRC expressions and implicit definitions of transformations, we considered inputs to be *free variables*. This is only a change in terminology, but it reflects the fact that in evaluating the interpretation on any input i_0 we will keep the interpretation

of the associated constant fixed, while we need to look at multiple bindings of the variables in each formula in order to form the output structure.

We will show that NRC expressions defining transformations from a nested relation of type T_1 to a nested relation of type T_2 correspond to a subset of interpretations of $\Sigma(T_2)$ within $\Sigma(T_1)$ that preserve $\mathfrak{U}$. The only additional restriction we impose is that all formulas $\varphi_{\mathsf{Domain}}^T$ and $\varphi_{\cong}^T$ in the definition of such an interpretation must be Δ_0. This forbids, for instance, universal quantification over the whole set of Ur-elements. We thus call a first-order interpretation of $\Sigma(T_2)$ within $\Sigma(T_1)$ consisting of Δ_0 formulas a Δ_0 *interpretation of $\Sigma(T_2)$ within $\Sigma(T_1)$*.

We now describe what it means for such an interpretation to define a transformation from an instance of one nested relational schema to another; that is, to map one object to another. We will denote the distinguished constant lying in the input sort by o_{in} and the distinguished constant in the output sort by o_{out}. Given any object o of type T, define M_o, a structure satisfying $\Sigma(T)$, as the least structure such that

- every subobject of o is part of M_o
- when $T_1 \times T_2$ is a subtype of T and a_1, a_2 are objects of sort T_1, T_2 of M_o, then $\langle a_1, a_2 \rangle$ is an object of M_o
- a copy of $\emptyset$ is part of M_o for every sort $\mathsf{Set}(T)$ in $\mathcal{SCH}_T$
- $\langle \rangle$ and $\{\langle \rangle\}$ are in M_o at sorts Unit and Bool.

The map $o \mapsto M_o$ shows how to translate an object to a logical structure that is appropriate as the input of an interpretation. Note that the further constraints ensure that every sort has at least one element in M_o and that there is one sort, $\mathfrak{U}$, which contains at least two elements; these technicality are important to ensure that interpretations are expressive enough.

We now discuss how the output of an interpretation is mapped back to an object. The output of an interpretation is a multi-sorted structure with a distinguished constant o_{out} encoding the output nested relational schema, but it is not technically a nested relational instance as required by our semantics for nested relational transformations. We can convert the output to a semantically appropriate entity via a modification of the well-known Mostowski collapse [Mos49]. We define $\mathsf{Collapse}(e, M)$ on elements e of the domain of a structure M for the multi-sorted encoding of a schema, by structural induction on the type of e:

- If e has sort $T_1 \times T_2$ then we set $\mathsf{Collapse}(e, M) = \langle \mathsf{Collapse}(\pi_1(e), M), \mathsf{Collapse}(\pi_2(e), M) \rangle$
- If e has sort $\mathsf{Set}(T)$, then we set $\mathsf{Collapse}(e, M) = \{\mathsf{Collapse}(t, M) \mid t \in e\}$
- Otherwise, if e has sort Unit or $\mathfrak{U}$, we set $\mathsf{Collapse}(e, M) = e$

We now formally describe how Δ_0 interpretations define functions between objects in the nested relational data model.

Definition 4.5. We say that a nested relational transformation $\mathcal{T}$ from T_1 to T_2 is defined by a Δ_0 interpretation $\mathcal{I}$ if, for every object o_{in} of type T_1, the structure M associated with o_{in} is mapped to M' where $\mathcal{T}(\mathsf{o}_{in})$ is equal to $\mathsf{Collapse}(\mathsf{o}_{out}, M')$.

We will often identify a Δ_0 interpretation with the corresponding transformation, speaking of its input and output as a nested relation (rather than the corresponding structure). For such an interpretation $\mathcal{I}$ and an input object o_{in} we write $\mathcal{I}(\mathsf{o}_{in})$ for the output of the transformation defined by $\mathcal{I}$ on o_{in}.

Example 4.6. Consider an input schema consisting of a single binary relation $R : \mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$, so an input object is a set of pairs, with each pair consisting of an Ur-element

and a set of Ur-elements. The corresponding theory is $\Sigma(\mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U})))$, which has sorts $T_{in} = \mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$, $\mathfrak{U} \times \mathsf{Set}(\mathfrak{U})$ and $\mathfrak{U}$ and relation symbols $\in_{\mathfrak{U}}$ and $\in_{\mathfrak{U} \times \mathsf{Set}(\mathfrak{U})}$ on top of equalities.

If we consider the following instance of the nested relational schema

$$R_0 = \{\langle a, \{a,b\}\rangle, \langle a, \{a,c\}\rangle, \langle b, \{a,c\}\rangle\}$$

Then the corresponding encoded structure M consists of:

- $M^{T_{in}}$ containing only the constant R_0
- $M^{\mathfrak{U} \times \mathsf{Set}(\mathfrak{U})}$ consisting of the elements of R_0,
- $M^{\mathfrak{U}}$ consisting of $\{a,b,c\}$
- $M^{\mathsf{Set}(\mathfrak{U})}$ consisting of the sets $\{a,b\}$, $\{a,c\}$,
- $M^{\mathsf{Unit}} = \{\langle\rangle\}$ and $M^{\mathsf{Bool}} = \{\emptyset, \{\langle\rangle\}\}$
- the element relations interpreted in the natural way

Consider the transformation that groups on the first component, returning an output object of type $O = \mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathsf{Set}(\mathfrak{U})))$. This is a variation of the grouping transformation from Example 3.4. On the example input R_0 the transformation would return

$$O_0 = \{\langle a, \{\{a,b\}, \{a,c\}\}\rangle, \langle b, \{\{a,c\}\}\rangle\}$$

The output would be represented by a structure having sorts $T_{out} = \mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathsf{Set}(\mathfrak{U})))$, $\mathfrak{U} \times \mathsf{Set}(\mathsf{Set}(\mathfrak{U}))$, $\mathfrak{U}$, $\mathsf{Set}(\mathsf{Set}(\mathfrak{U}))$ and $\mathsf{Set}(\mathfrak{U})$ in addition to Unit and Bool. It is easy to capture this transformation with a Δ_0 interpretation. For example, the interpretation could code the output sort $\mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathsf{Set}(\mathfrak{U})))$ by the sort $\mathsf{Set}(\mathfrak{U} \times \mathsf{Set}(\mathfrak{U}))$, representing each group by the corresponding Ur-element.

We will often make use of the following observation about interpretations:

Proposition 4.7. Δ_0 *interpretations can be composed, and their composition correspond to the underlying composition of transformations.*

The composition of nested relational interpretations amounts to the usual composition of FO-interpretations (see e.g. [BK09]) and an easy check that the additional requirements we impose on nested relational interpretations are preserved.

We can now state the equivalence of NRC and interpretations formally:

Theorem 4.8. *Every transformation in* NRC *can be translated effectively to a Δ_0 interpretation. Conversely, for every Δ_0 interpretation, one can effectively form an equivalent* NRC *expression. The translation from* NRC *to interpretations can be done in* EXPTIME *while the converse translation can be performed in* PTIME.

This characterization holds when equivalence is over finite nested relational inputs and also when arbitrary nested relations are allowed as inputs to the transformations.

Note that very similar results occur in the literature, going back at least to [Van01]. Thus we defer the proof to Appendix E. The direction from interpretations to NRC will be the one that is directly relevant to us in the sequel, and its proof is given by a simple translation.

4.3. **Reduction to a characterization theorem in multi-sorted logic.** The first step in the proof of Theorem 4.2 is to reduce to a more general statement relating implicit definitions in multi-sorted logic to interpretations.

The first part of the reduction is to argue that we can suppress auxiliary parameters $\vec{a}$ in implicit definitions, proving the equivalence of the first two items in 4.3:

Lemma 4.9. *For any Δ_0 formula $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a})$ that implicitly defines o_{out} as a function of o_{in}, there is another Δ_0 formula $\Sigma'(\mathsf{o}_{in}, \mathsf{o}_{out})$ which implicitly defines o_{out} as a function of o_{in}, such that $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a}) \to \Sigma'(\mathsf{o}_{in}, \mathsf{o}_{out})$.*

Note that from this lemma we get the equivalence of 1 and 2 in Corollary 4.3.

Proof. The assumption that Σ implicitly defines o_{out} as a function of o_{in} means that we have an entailment

$$\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a}) \models \Sigma(\mathsf{o}_{in}, \mathsf{o}'_{out}, \vec{a'}) \to \mathsf{o}_{out} = \mathsf{o}'_{out}$$

Applying Δ_0 interpolation, Proposition 4.4, we may obtain a formula $\theta(\mathsf{o}_{in}, \mathsf{o}_{out})$ such that

$$\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a}) \models \theta(\mathsf{o}_{in}, \mathsf{o}_{out}) \qquad \text{and} \qquad \theta(\mathsf{o}_{in}, \mathsf{o}_{out}) \wedge \Sigma(\mathsf{o}_{in}, \mathsf{o}'_{out}, \vec{a'}) \models \mathsf{o}_{out} = \mathsf{o}'_{out}$$

Now we can derive the following entailment

$$\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a}) \models [\theta(\mathsf{o}_{in}, \mathsf{o}'_{out}) \wedge \theta(\mathsf{o}_{in}, \mathsf{o}''_{out})] \to \mathsf{o}'_{out} = \mathsf{o}''_{out}$$

This entailment is obtained from the second property of θ, since we can infer that $\mathsf{o}'_{out} = \mathsf{o}_{out}$ and $\mathsf{o}''_{out} = \mathsf{o}_{out}$.

Now we can apply interpolation again to obtain a formula $D(\mathsf{o}_{in})$ such that

$$\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a}) \models D(\mathsf{o}_{in}) \qquad \text{and} \qquad D(\mathsf{o}_{in}) \wedge \theta(\mathsf{o}_{in}, \mathsf{o}'_{out}) \wedge \theta(\mathsf{o}_{in}, \mathsf{o}''_{out}) \models \mathsf{o}'_{out} = \mathsf{o}''_{out}$$

We now claim that $\Sigma'(\mathsf{o}_{in}, \mathsf{o}_{out}) := D(\mathsf{o}_{in}) \wedge \theta(\mathsf{o}_{in}, \mathsf{o}_{out})$ is an implicit definition extending Σ. Functionality of Σ' is a consequence of the second entailment witnessing that D is an interpolant. Finally, the implication $\exists \vec{a} \, \Sigma(\mathsf{o}_{in}, \mathsf{o}_{out}, \vec{a}) \models \Sigma'(\mathsf{o}_{in}, \mathsf{o}_{out})$ is given by the combination of the first entailments witnessing that θ and D are interpolants. $\qquad \square$

With the above result in hand, from this point on we assume that we do not have auxiliary parameters $\vec{a}$ in our implicit definitions.

Reduction to Monadic schemas. A *monadic type* is a type built only using the atomic type $\mathfrak{U}$ and the type constructor Set. To simplify notation we define $\mathfrak{U}_0 := \mathfrak{U}$, $\mathfrak{U}_1 := \mathsf{Set}(\mathfrak{U}_0), \ldots, \mathfrak{U}_{n+1} := \mathsf{Set}(\mathfrak{U}_n)$. A monadic type is thus a $\mathfrak{U}_n$ for some $n \in \mathbb{N}$. A nested relational schema is monadic if it contains only monadic types, and a Δ_0 formula is said to be monadic if all of its variables have monadic types.

Restricting to monadic formulas simplifies our arguments considerably. It turns out that by the usual "Kuratowski encoding" of pairs by sets, we can reduce all of our questions about implicit versus explicit definability to the case of monadic schemas. The following proposition implies that we can derive all of our main results for arbitrary schemas from their restriction to monadic formulas. The proof is routine but tedious, so we defer it to Appendix C–D.

Proposition 4.10. *For any nested relational schema $\mathcal{SCH}$, there is a monadic nested relational schema $\mathcal{SCH}'$, an injection $\mathsf{Convert}$ from instances of $\mathcal{SCH}$ to instances of $\mathcal{SCH}'$ that is definable in NRC, and an $\mathsf{NRC}[\mathrm{GET}]$ expression $\mathsf{Convert}^{-1}$ such that $\mathsf{Convert}^{-1} \circ \mathsf{Convert}$ is the identity transformation from $\mathcal{SCH} \to \mathcal{SCH}$.*

Furthermore, there is a Δ_0 formula $\mathsf{Im}_{\mathsf{Convert}}$ from $\mathcal{SCH}'$ to Bool such that $\mathsf{Im}_{\mathsf{Convert}}(i')$ holds if and only if $i' = \mathsf{Convert}(i)$ for some instance i of $\mathcal{SCH}$.

These translations can also be given in terms of Δ_0 interpretations rather than NRC expressions.

As we now explain, Proposition 4.10 allows us to reduce Theorem 4.2 to the special case where we have only monadic nested relational schemas. Given a Δ_0 implicit definition $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out})$ we can form a new definition that computes the composition of the following transformations: $\mathsf{Convert}^{-1}_{\mathcal{SCH}_{in}}$, a projection onto the first component, the transformation defined by Σ, and $\mathsf{Convert}_{\mathcal{SCH}_{out}}$. Our new definition captures this composition by a formula $\Sigma'(\mathsf{o}'_{in}, \mathsf{o}'_{out})$ that defines o'_{out} as a function of o_{in}, where the formula is over a monadic schema. Assuming that we have proven the theorem in the monadic case, we would get an NRC expression E' from $\mathcal{SCH}'_{in}$ to $\mathcal{SCH}'_{out}$ agreeing with this formula on its domain. Now we can compose $\mathsf{Convert}_{\mathcal{SCH}_{in}}$, E', $\mathsf{Convert}^{-1}_{\mathcal{SCH}_{out}}$, and the projection to get an NRC expression agreeing with the partial function defined by $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out})$ on its domain, as required.

Reduction to a result in multi-sorted logic. Now we are ready to give our last reduction, relating Theorem 4.2 to a general result concerning multi-sorted logic.

Let $\mathcal{SIG}$ be any multi-sorted signature, Sorts_1 be its sorts and Sorts_0 be a subset of Sorts_1. We say that a relation R is *over* Sorts_0 if all of its arguments are in Sorts_0. Let Σ be a set of sentences in $\mathcal{SIG}$. Given a model M for $\mathcal{SIG}$, let $\mathsf{Sorts}_0(M)$ be the union of the domains of relations over Sorts_0, and let $\mathsf{Sorts}_1(M)$ be defined similarly.

We say that Sorts_1 is *implicitly interpretable* over Sorts_0 relative to Σ if:

> *Fix any models M_1 and M_2 of Σ. Suppose m is an isomorphism from $\mathsf{Sorts}_0(M_1)$ to $\mathsf{Sorts}_0(M_2)$: that is a bijection from the domain of each sort, that preserves all relations over Sorts_0 in both directions. Then m extends to a unique mapping from $\mathsf{Sorts}_1(M_1)$ to $\mathsf{Sorts}_1(M_2)$ which preserves all relations over Sorts_1.*

Informally, implicit interpretability states that the sorts in Sorts_1 are semantically determined by the sorts in Sorts_0. The property implies in particular that if M_1 and M_2 agree on the interpretation of sorts in Sorts_0, then the identity mapping on sorts in Sorts_0 extends to a mapping that preserves sorts in Sorts_1.

We relate this semantic property to a syntactic one. We say that Sorts_1 is *explicitly interpretable* over Sorts_0 relative to Σ if for all S in Sorts_1 there is a formula $\psi_{\mathsf{S}}(\vec{x}, y)$ where $\vec{x}$ are variables with sorts in Sorts_0, y a variable of sort S, such that:

- In any model M of Σ, ψ_{S} defines a partial function F_{S} mapping Sorts_0 tuples surjectively on to S.
- For every relation R of arity n over Sorts_1, there is a formula $\psi_R(\vec{x}_1, \ldots \vec{x}_n)$ using only relations of Sorts_0 and only quantification over Sorts_0 such that in any model M of Σ, the pre-image of R under the mappings F_{S} for the different arguments of R is defined by $\psi_R(\vec{x}_1, \ldots \vec{x}_n)$.

Explicit interpretability states that there is an interpretation in the sense of the previous section that produces the structure in Sorts_1 from the structure in Sorts_0, and in addition there is a definable relationship between an element e of a sort in Sorts_1 and the tuple that codes e in the interpretation. Note that ψ_{S}, the mapping between the elements y in S and the tuples in Sorts_0 that interpret them, can use arbitrary relations. The key property is that when we pull a relation R over Sorts_1 back using the mappings ψ_{S}, then we obtain something definable using Sorts_0.

With these definitions in hand, we are ready to state a result in multi-sorted logic which allows us to generate interpretations from classical proofs of functionality:

Theorem 4.11. *For any $\Sigma, \mathsf{Sorts}_0, \mathsf{Sorts}_1$ such that Σ entails that a sort of Sorts_0 has at least two elements, Sorts_1 is explicitly interpretable over Sorts_0 if and only if it is implicitly interpretable over Sorts_0.*

This can be thought of as an analog of Beth's theorem [Bet53, Cra57b] for multi-sorted logic. The proof is given in the next subsection. For now we explain how it implies Theorem 4.2. In this explanation we assume a monadic schema for both input and output. Thus every element e in an instance has sort $\mathfrak{U}_n$ for some $n \in \mathbb{N}$.

Consider a Δ_0 formula $\Sigma(\mathsf{o}_{in}, \mathsf{o}_{out})$ over a monadic schema that implicitly defines o_{out} as a function of o_{in}. Σ can be considered as a multi-sorted first-order formula with sorts for every subtype occurrence of the input as well as distinct sorts for every subtype occurrence of the output other than $\mathfrak{U}$. Because we are dealing with monadic input and output schema, every sort other than $\mathfrak{U}$ will be of the form $\mathsf{Set}(T)$, and these sorts have only the element relations $\in_T$ connecting them. We refer to these as *input sorts* and *output sorts*. We modify Σ by asserting that all elements of the input sorts lie underneath o_{in}, and all elements of the output sorts lie underneath o_{out}. Since Σ was Δ_0, this does not change the semantics. We also conjoin to Σ the sanity axioms for the schema, including the extensionality axiom at the sorts corresponding to each object type. Let Σ^* be the resulting formula. In this transformation, as was the case with interpretations, we change our perspective on inputs and outputs, considering them as constants rather than as free variables. We do this only to match our result in multi-sorted logic, which deals with a set of *sentences* in multi-sorted first-order logic, rather than formulas with free variables.

Given models M and M' of Σ^*, we define relations $\equiv_i$ connecting elements of M of depth i with elements of M' of depth i. For $i = 0$, $\equiv_i$ is the identity: that is, it connects elements of $\mathfrak{U}$ if and only if they are identical. For $i = j + 1$, $\equiv_i (x, x')$ holds exactly when for every $y \in x$ there is $y' \in x'$ such that $y \equiv_j y'$, and vice versa.

The fact that Σ implicitly defines o_{out} as a function of o_{in} tells us that:

> *Suppose $M \models \Sigma^*$, $M' \models \Sigma^*$ and M and M' are identical on the input sorts. Then the mapping taking a $y \in M$ of depth i to a $y' \in M'$ such that $y' \equiv_i y$ is an isomorphism of the output sorts that is the identity on $\mathfrak{U}$. Further, any isomorphism of $\mathsf{Sorts}_1(M)$ on to $\mathsf{Sorts}_1(M')$ that is the identity on $\mathfrak{U}$ must be equal to M: one can show this by induction on the depth i using the fact that Σ^* includes the extensionality axiom.*

From this, we see that the output sorts are implicitly interpretable over the input sorts relative to Σ^*. Using Theorem 4.11, we conclude that the output sorts are explicitly interpretable in the input sorts relative to Σ^*. Applying the conclusion to the formula $x = x$, where x is a variable of a sort corresponding to object type T of the output, we obtain a first-order formula $\varphi^T_{\mathsf{Domain}}(\vec{x})$ over the input sorts. Applying the conclusion to the formula $x = y$ for x, y variables corresponding to the object type T we get a formula $\varphi_{\equiv_T}(\vec{x}, \vec{x}')$ over the input sorts. Finally applying the conclusion to the element relation ϵ_T at every level of the output, we get a first-order formula $\varphi_{\epsilon_T}(\vec{x}, \vec{x}')$ over the input sorts. Because Σ^* asserts that each element of the input sorts lies beneath a constant for o_{in}, we can convert all quantifiers to bind only beneath o_{in}, giving us Δ_0 formulas. It is easy to verify that these formulas give us the desired interpretation. This completes the proof of Theorem 4.2, assuming Theorem 4.11.

4.4. **Proof of the multi-sorted logic result.** In the previous subsection we reduced our goal result about generating interpretations from proofs to a result in multi-sorted first-order logic, Theorem 4.11. We will now present the proof of Theorem 4.11. The direction from explicit interpretability to implicit interpretability is straightforward, so we will be interested only in the direction from implicit to explicit. Although the theorem appears to be new as stated, each of the components is a variant of arguments that already appear in the model theory literature. The core of our presentation here is a variation of the proof of Gaifman's coordinatisation theorem as presented in [Hod93].

We make use of only quite basic results from model theory:

- the *compactness theorem* for first-order logic, which states that for any theory Γ, if every finite subcollection of Γ is satisfiable, then Γ is satisfiable;
- the *downward Löwenheim-Skolem theorem*, which states that if Γ is countable and has a model, then it has a countable model;
- the *omitting types theorem* for first-order logic. A first-order theory Σ is said to be *complete* if for every other first-order sentence φ in the vocabulary of Σ, either φ or $\neg\varphi$ is entailed by Σ. Given a set of constants B, a *type* over B is an infinite collection $\tau(\vec{x})$ of formulas using variables $\vec{x}$ and constants B. A type is *complete* with respect to a theory Σ if every first-order formula with variables in $\vec{x}$ and constants from B is either entailed or contradicted by $\tau(\vec{x})$ and Σ. A type τ is said to be *realized* in a model M if there is a $\vec{x}_0$ in M satisfying all formulas in τ. τ is *non-principal* (with respect to a first-order theory Σ) if there is no formula $\gamma_0(\vec{x})$ such that $\Sigma \wedge \gamma_0(\vec{x})$ entails all of $\tau(\vec{x})$. The version of the omitting types theorem that we will use states that

  > *if we have a countable set Γ of complete types that are all non-principal relative to a complete theory Σ, there is some model M of Σ in which none of the types in Γ are realized.*

Each of these results follows from a standard model construction technique [Hod93].

We can easily show that to prove the multi-sorted result, it suffices to consider only those Σ which are complete theories.

Proposition 4.12. *Theorem 4.11 follows from its restriction to Σ a complete theory.*

Recall that we are proving the direction from implicit interpretability to explicit interpretability. Our first step will be to show that each element in the output sort is definable from the input sorts, if we allow ourselves to guess some parameters. For example, consider the grouping transformation mentioned in Example 3.4. Each output is obtained from grouping input relation F over some Ur-element a. So each member of the output is definable from the input constant F and a "guessed" input element a. We will show that this is true in general.

For the next results, we have a *blanket assumption that our underlying language is countable*, which will be necessary in some applications of the compactness theorem.

Given a model M of Σ and $\vec{x}_0 \in \mathsf{Sorts}_1$ within M, the *type of $\vec{x}_0$ with parameters from Sorts_0* is the set of all formulas satisfied by $\vec{x}_0$, using any sorts and relations but only constants from Sorts_0.

A type p is *isolated over* Sorts_0 if there is a formula $\varphi(\vec{x}, \vec{a})$ with parameters $\vec{a}$ from Sorts_0 such that $M \models \varphi(\vec{x}, \vec{a}) \to \gamma(\vec{x})$ for each $\gamma \in p$. The following is a step towards showing that elements in the output are well-behaved:

Lemma 4.13. *Suppose* Sorts_1 *is implicitly interpretable over* S_0 *with respect to* Σ. *Then in any countable model* M *of* Σ *the type of any* $\vec{b}$ *over* Sorts_1 *with parameters from* Sorts_0 *is isolated over* Sorts_0.

Proof. Fix a counterexample $\vec{b}$, and let Γ be the set of formulas in Sorts_1 with constants from Sorts_0 satisfied by $\vec{b}$ in M. We claim that there is a model M' with $\mathsf{Sorts}_0(M')$ identical to $\mathsf{Sorts}_0(M)$ where there is no tuple satisfying Γ. This follows from the failure of isolation and the omitting types theorem.

Now we have a contradiction of implicit interpretability, since the identity mapping on Sorts_0 can not extend to an isomorphism of relations over Sorts_1 from M to M'. $\square$

The next step is to argue that every element of Sorts_1 is definable by a formula using parameters from Sorts_0.

Lemma 4.14. *Assume implicit interpretability of* Sorts_1 *over* Sorts_0 *relative to* Σ. *In any model* M *of* Σ, *for every element* e *of a sort* S_1 *in* Sorts_1, *there is a first-order formula* $\psi_e(\vec{y}, x)$ *with variables* $\vec{y}$ *having sort in* Sorts_0 *and* x *a variable of sort* S_1, *along with a tuple* $\vec{a}$ *in* $\mathsf{Sorts}_0(M)$ *such that* $\psi_e(\vec{a}, x)$ *is satisfied only by* e *in* M.

In a single-sorted setting, this can be found in [Hod93] Theorem 12.5.8 where it is referred to as "Gaifman's coordinatisation theorem", credited independently to unpublished work of Haim Gaifman and Dale Myers. The multi-sorted version is also a variant of Remark 1.2, part 4 in [Hru14], which points to a proof in the appendix of [CH99]; the remark assumes that Sorts_1 is the set of all sorts. Another variation is Theorem 3.3.4 of [AMN08].

Proof. Since a counterexample involves only formulas in a countable language, by the Löwenheim-Skolem theorem mentioned above, it is enough to consider the case where M is countable. By Lemma 4.13, the type of every e is isolated by a formula $\varphi(\vec{x}, \vec{a})$ with parameters from Sorts_0 and relations from Sorts_1. We claim that φ defines e: that is, e is the only satisfier. If not, then there is $e' \neq e$ that satisfies φ. Consider the relation $\vec{j} \equiv \vec{j}'$ holding if $\vec{j}$ and $\vec{j}'$ satisfy all the same formulas using relations and variables from Sorts_1 and parameters from Sorts_0. Isolation implies that $e \equiv e'$. Further, isolation of types shows that $\equiv$ has the "back-and-forth property" given $\vec{d} \equiv \vec{d}'$, and $\vec{c}$ we can obtain $\vec{c}'$ with $\vec{d}\vec{c} \equiv \vec{d}'\vec{c}'$. To see this, fix $\vec{d} \equiv \vec{d}'$ and consider $\vec{c}$. We have $\gamma(\vec{x}, \vec{y}, \vec{a})$ isolating the type of $\vec{d}, \vec{c}$, and further $\vec{d}$ satisfies $\exists \vec{y}\, \gamma(\vec{x}, \vec{y}, \vec{a})$ and thus so does $\vec{d}'$ with witness $\vec{c}'$. But then using $\vec{d} \equiv \vec{d}'$ again we see that $\vec{d}, \vec{c} \equiv \vec{d}', \vec{c}'$. Using countability of M and this property we can inductively create a mapping on M fixing Sorts_0 pointwise, preserving all relations in Sorts_1, and taking e to e'. But this contradicts implicit interpretability. $\square$

Lemma 4.15. *The formula in Lemma 4.14 can be taken to depend only on the sort* S_1.

Proof. Consider the type over the single variable x in S_1 consisting of the formulas $\neg \delta_\varphi(x)$ where $\delta_\varphi(x)$ is the following formula

$$\exists \vec{b}\, [\varphi(\vec{b}, x) \wedge \forall x'\, (\varphi(\vec{b}, x') \rightarrow x' = x)]$$

where the tuple $\vec{b}$ ranges over Sorts_0. By Lemma 4.14, this type cannot be satisfied in a model of Σ. Since it is unsatisfiable, by compactness, there are finitely many formulas $\varphi_1(\vec{b}, x), \ldots, \varphi_n(\vec{b}, x)$ such that $\forall x\, \bigvee_{i=1}^{n} \delta_{\varphi_i}(x)$ is satisfied. Therefore, each $\varphi_i(\vec{b}, x)$ defines a partial function from tuples of S_0 to S_1 and every element of S_1 is covered by one of the φ_i. Recall that we assumed that Σ enforces that Sorts_0 has a sort with at least two elements.

Thus we can combine the $\varphi_i(\vec{b}, x)$ into a single formula $\psi(\vec{b}, \vec{c}, x)$ defining a surjective partial function from S_0 to S_1 where $\vec{c}$ is an additional parameter in Sorts_0 selecting some $i \leq n$. $\quad\square$

We now need to go from the "sub-definability" or "element-wise definability" result above to an interpretation. Consider the formulas ψ_S produced by Lemma 4.15. For a relation R of arity n over Sorts_1, where the i^{th} argument has sort S_i, consider the formula

$$\psi_R(\vec{y}_1 \ldots \vec{y}_n) = \exists x_1 \ldots x_n \; [R(x_1 \ldots x_n) \wedge \bigwedge_i \psi_{\mathsf{S}_i}(\vec{y}_i, x_i)]$$

where $\vec{x}_i$ is a tuple of variables of sorts in Sorts_0.

The formulas ψ_S for each sort S and the formulas ψ_R for each relation R are as required by the definition of explicitly interpretable, except that they may use quantified variables and relations of Sorts_1, while we only want to use variables and relations from Sorts_0. We take care of this in the following lemma, which says that formulas over Sorts_1 do not allow us to define any more subsets of Sorts_0 than we can with formulas over Sorts_0.

Lemma 4.16. *Under the assumption of implicit interpretability, for every formula $\varphi(\vec{y})$ over Sorts_1 with $\vec{y}$ variables whose sorts are in Sorts_0 there is a formula $\varphi^\circ(\vec{y})$ over Sorts_0 – that is, containing only variables, constants, and relations from Sorts_0 – such that for every model M of Σ,*

$$M \models \forall \vec{y} \; [\varphi(\vec{y}) \leftrightarrow \varphi^\circ(\vec{y})]$$

Proof. We give an argument assuming the existence of a saturated model for the theory: that is, a model M in which for every set of formulas $\Gamma(\vec{x})$, with parameters from the model, of cardinality strictly smaller than the model, if Γ is finitely satisfiable in M then it is realized in M. Such models exist for any theory under the generalized continuum hypothesis GCH. The additional set-theoretic hypothesis can be removed by using weaker notions of saturation – the modification is a standard one in model theory, see [CK92, Hod93].

Assume not, with $\varphi(\vec{y})$ as a counterexample. By completeness and our assumption, we know that there is a saturated model M of Σ containing $\vec{c}, \vec{c}'$ that agree on all formulas in Sorts_0 but that disagree on φ. Call R the reduct of M to Sorts_0. The partial map that sends $\vec{c}$ to $\vec{c}'$ is a partial automorphism of R. Since M is saturated, so is R, so in particular R is strongly homogeneous and the aforementioned partial map can be extended to an automorphism of R that sends $\vec{c}$ to $\vec{c}'$. But then by weak implicit interpretability, this should extend to an automorphism of M that sends $\vec{c}$ to $\vec{c}'$. This implies that M satisfies $\varphi(\vec{c}) \leftrightarrow \varphi(\vec{c}')$, a contradiction. $\quad\square$

Above we obtained the formulas ψ_R for each relation symbol R needed for an explicit interpretation. We can obtain formulas defining the necessary equivalence relations $\psi_\equiv$ and ψ_{Domain} easily from these.

Putting Lemmas 4.14, 4.15, and 4.16 together yields a proof of Theorem 4.11.

4.5. **Putting it all together to complete the proof of Theorem 4.2.** We summarize our results on extracting NRC expressions from classical proofs of functionality. We have shown in Subsection 4.3 how to convert the problem to one with no extra variables other than input and output and with only monadic schemas – and thus no use of products or tupling. We also showed how to convert the resulting formula into a theory in multi-sorted first-order logic. That is, we no longer need to talk about Δ_0 formulas.

In Subsection 4.4 we showed that from a theory in multi-sorted first-order logic we can obtain an interpretation. This first-order interpretation in a multi-sorted logic can then be converted back to a Δ_0 interpretation, since the background theory forces each of the input sorts in the multi-sorted structure to correspond to a level of nesting below one of the constants corresponding to an input object. Finally, the results of Subsection 4.2 allow us to convert this interpretation to an NRC expression. With the exception of the result in multi-sorted logic, all of the constructions are effective. Further, these effective conversions are all in polynomial time except for the transformation from an interpretation to an NRC expression, which is exponential time in the worst case. Outside of the multi-sorted result, which makes use of infinitary methods, the conversions are each sound when equivalence over finite input structures is considered as well as the default case when arbitrary inputs are considered. As explained in Subsection 4.3, when equivalence over finite inputs is considered, we cannot hope to get a synthesis result of this kind.

5. The effective result: efficiently generating NRC expressions from proofs

We now turn to the question of *effective* conversion from implicit definitions to explicit NRC transformations, leading up to our second main contribution.

5.1. Moving effectively from implicit to explicit: statement of the main result.

Recall that previously we have phased implicit definability as an entailment in the presence of extensionality axioms. We also recall that we can rephrase it without explicitly referring to extensionality.

A Δ_0 formula $\varphi(\vec{i}, \vec{a}, o)$, implicitly defines variable o in terms of variables $\vec{i}$ if we have

$$\varphi(\vec{i}, \vec{a}, o) \ \wedge \ \varphi(\vec{i}, \vec{a}', o') \ \models_{\mathsf{nested}} \ o \equiv_T o' \tag{5.1}$$

Here $\equiv_T$ is equivalence-modulo-extensionality, as defined in Section 3. It is the same as equality if we add extensionality axioms on the left of the entailment symbol. Thus *this entailment is the same as entailment with the sanity axioms from the previous section, using native equality rather than $\equiv_T$.* And since these are Δ_0 formulas, *it is also the same as entailment over "general models", where we assume only correct typing and projection commuting with tupling.*

Proof systems for Δ_0 formulas. Recall that our goal is an effective version of Corollary 4.3. For this we need to formalize our proof system for Δ_0 formulas, which will allow us to talk about proof witnesses for implicit definability.

A finite multiset of primitive membership expressions $t \in u$ (i.e. extended Δ_0 formulas) will be called an $\in$-*context*. These expressions arise naturally when breaking down bounded quantifiers during a proof.

We introduce notation for instantiating a block of bounded quantifiers at a time in a Δ_0 formula.

A term is a *tuple-term* if it is built up from variables using pairing.

If we want to talk only about *effective generation* of NRC witnesses from proofs, we can use a basic proof system for Δ_0 formulas, whose inference rules are shown in Figure 2.

The node labels are a variation of the traditional rules for first-order logic, with a couple of quirks related to the specifics of Δ_0 formulas. Each node label has shape $\Theta; \Gamma \vdash \Delta$ where

$$\text{Ax}\ \frac{}{\Theta;\ \Gamma, \varphi \vdash \varphi, \Delta} \qquad\qquad \bot\text{-L}\ \frac{}{\Theta;\ \Gamma, \bot \vdash \Delta}$$

$$\neg\text{-L}\ \frac{\Theta;\ \Gamma \vdash \neg\varphi, \Delta}{\Theta;\ \Gamma, \varphi \vdash \Delta} \qquad\qquad \neg\text{-R}\ \frac{\Theta;\ \Gamma, \varphi \vdash \Delta}{\Theta;\ \Gamma \vdash \neg\varphi, \Delta}$$

$$\wedge\text{-R}\ \frac{\Theta;\ \Gamma \vdash \varphi_1, \Delta \qquad \Theta;\ \Gamma \vdash \varphi_2, \Delta}{\Theta;\ \Gamma \vdash \Delta, \varphi_1 \wedge \varphi_2} \qquad\qquad \vee\text{-R}\ \frac{\Theta;\ \Gamma \vdash \varphi_1, \varphi_2, \Delta}{\Theta;\ \Gamma \vdash \varphi_1 \vee \varphi_2, \Delta}$$

$$\forall\text{-R}\ \frac{\Theta, y \in b;\ \Gamma \vdash \varphi[y/x], \Delta}{\Theta;\ \Gamma \vdash \forall x \in b\ \varphi, \Delta}\ y\ \text{FRESH} \qquad\qquad \exists\text{-R}\ \frac{\Theta, t \in b;\ \Gamma \vdash \varphi[t/x], \exists x \in b\ \varphi, \Delta}{\Theta, t \in b;\ \Gamma \vdash \exists x \in b\ \varphi, \Delta}$$

$$\text{REFL}\ \frac{\Theta;\ \Gamma, t =_{\mathfrak{U}} t \vdash \Delta}{\Theta;\ \Gamma \vdash \Delta} \qquad\qquad \text{REPL}\ \frac{\Theta;\ \Gamma, t =_{\mathfrak{U}} u, \varphi[u/x], \varphi[t/x] \vdash \Delta}{\Theta;\ \Gamma, t =_{\mathfrak{U}} u, \varphi[t/x] \vdash \Delta}$$

$$\times_\eta\ \frac{\Theta[\langle x_1, x_2\rangle/x];\ \Gamma[\langle x_1, x_2\rangle/x] \vdash \Delta[\langle x_1, x_2\rangle/x]}{\Theta;\ \Gamma \vdash \Delta}\ x_1, x_2\ \text{FRESH}$$

$$\times_\beta\ \frac{\Theta[t_i/x];\ \Gamma[t_i/x] \vdash \Delta[t_i/x]}{\Theta[\pi_i(\langle t_1, t_2\rangle)/x];\ \Gamma[\pi_i(\langle t_1, t_2\rangle)/x] \vdash \Delta[\pi_i(\langle t_1, t_2\rangle)/x]}\ i \in \{1, 2\}$$

Figure 2: Proof rules for a Δ_0 calculus, without restrictions for efficient generation of witnesses. The left side of ; specifies the $\in$-context. The negation rules $\neg$-L and $\neg$-R permit to exchange formulas between both sides of $\vdash$ such that it suffices to have the rules for the connectives only for one side, where we choose the right side.

- Θ is an $\in$-context. Recall that these are multisets of membership atoms — the only formulas in our proof system that are extended Δ_0 but not Δ_0. They will emerge during proofs involving Δ_0 formulas when we start breaking down bounded-quantifier formulas.
- Γ and Δ are finite multisets of Δ_0 formulas.[1]

For example, REPL in the figure is a "congruence rule", capturing that terms that are equal are interchangeable. Informally, it says that to prove conclusion Δ from a hypothesis that includes a formula φ including variable t and an equality $t =_{\mathfrak{U}} u$, it suffices to add to the hypotheses a copy of φ with u replacing some occurrences of t.

A proof tree whose root is labelled by $\Theta;\ \Gamma \vdash \Delta$ witnesses that, for any given meaning of the free variables, if all the membership relations in Θ and all formulas in Γ are satisfied, then there is a formula in Δ which is true. We say that we have a proof of a single formula φ when we have a proof of $\emptyset; \emptyset \vdash \varphi$.

The proof system is easily seen to be sound: if $\Theta;\ \Gamma \vdash \Delta$, then $\Theta;\ \Gamma \models \Delta$, where we remind the reader that $\models$ considers all models, not just extensional ones. It can be shown to be complete by a standard technique (a "Henkin construction", see Appendix F).

---

[1] Much of our machinery also works if sequents are built from finite sets instead of finite multisets. However, the specification of certain proof transformations (e.g. Appendix G) is much simpler with multisets.

$$= \frac{}{\Theta \vdash x =_{\mathfrak{U}} x, \Delta} \qquad\qquad \top \frac{}{\Theta \vdash \top, \Delta}$$

$$\neq \frac{\Theta \vdash t \neq_{\mathfrak{U}} u, \alpha[u/x], \alpha[t/x], \Delta^{\mathsf{EL}}}{\Theta \vdash t \neq_{\mathfrak{U}} u, \alpha[t/x], \Delta^{\mathsf{EL}}} \; \alpha \; \text{ATOMIC}$$

$$\wedge \frac{\Theta \vdash \varphi_1, \Delta \qquad \Theta \vdash \varphi_2, \Delta}{\Theta \vdash \varphi_1 \wedge \varphi_2, \Delta} \qquad\qquad \vee \frac{\Theta \vdash \varphi_1, \varphi_2, \Delta}{\Theta \vdash \varphi_1 \vee \varphi_2, \Delta}$$

$$\forall \frac{\Theta, y \in b \vdash \varphi[y/x], \Delta}{\Theta \vdash \forall x \in b. \; \varphi, \Delta} \; y \; \text{FRESH} \qquad \exists \frac{\Theta, t \in b; \; \Gamma \vdash \varphi[t/x], \exists x \in b \; \varphi, \Delta^{\mathsf{EL}}}{\Theta, t \in b; \Gamma \vdash \exists x \in b \; \varphi, \Delta^{\mathsf{EL}}} \; t \; \text{TUPLE-TERM}$$

$$\times_\eta \frac{\Theta[\langle x_1, x_2 \rangle / x] \vdash \Delta^{\mathsf{EL}}[\langle x_1, x_2 \rangle / x]}{\Theta \vdash \Delta^{\mathsf{EL}}} \; x_1, x_2 \; \text{FRESH}$$

$$\times_\beta \frac{\Theta[t_i/x] \vdash \Delta^{\mathsf{EL}}[t_i/x]}{\Theta[\pi_i(\langle t_1, t_2 \rangle)/x] \vdash \Delta^{\mathsf{EL}}[\pi_i(\langle t_1, t_2 \rangle)/x]} \; i \in \{1, 2\}$$

Figure 3: Our EL-normalized calculus for efficient generation of witnesses. The left side of $\vdash$ specifies the $\in$-context. The right side a finite multiset of Δ_0 formulas. Recall that an atomic formula is an equality or inequality for $\mathfrak{U}$. Formula multisets Δ^{EL} are existential-leading, that is, they contain only atomic formulas and formulas with existential quantification as top-level connective.

To generate NRC definitions *efficiently* from proof witnesses will require a more restrictive proof system, in which we enforce some ordering on how proof rules can be applied, depending on the shape of the hypotheses. We refer to proofs in this system as EL-*normalized proofs*, To this end we consider a multiset of formulas *existential-leading* (EL) if it contains only atomic formulas (i.e. formulas of the form $t =_{\mathfrak{U}} t'$ or $t \neq_{\mathfrak{U}} t'$), formulas with existential quantification as top-level connective and the truth-value constant $\bot$.

Our restricted proof system – EL-normalized proofs – is shown in Figure 3. A superficial difference from Figure 2 is that the restricted system is "almost 1-sided": Δ_0 formulas only occur on the right, with only $\in$-contexts on the left. In particular, a top-level goal $\Theta; \Gamma \vdash \Delta$ in the higher-level system would be expressed as $\Theta \vdash \neg\Gamma, \Delta$ in this system. We will often abuse notation by referring to EL-normalized proofs of a 2-sided sequent $\Theta; \Gamma \vdash \Delta$, considering them as "macros" for the corresponding 1-sided sequent. For example, the hypothesis of the $\neq$ rule could be written in 2-sided notation as $\Theta, t =_{\mathfrak{U}} u \vdash \alpha[u/x], \alpha[t/x], \Delta^{\mathsf{EL}}$ while the conclusion could be written as $\Theta, t =_{\mathfrak{U}} u \vdash \alpha[t/x], \Delta^{\mathsf{EL}}$. As with REPL in the prior system, this rule is about duplicating a hypothesis with some occurrences of t replaced by u.

A major aspect of the restriction is baked into the shape of the $\exists$ rule. It enforces that existentials are instantiated only in a context which is EL, and that the term being substituted is of a simple shape (tupling of variables).

Soundness is evident, since it is a special case of the proof system above. Completeness is not as obvious, since we are restricting the proof rules. But we can translate proofs in

the more general system of Figure 2 into a EL-normalized proof, but with an exponential blow-up: see Appendix G for details.

Furthermore, since for Δ_0 formulas equivalence over all structures is the same as equivalence over nested relations, a Δ_0 formula φ is provable exactly when $\models_{\mathsf{nested}} \varphi$.

Example 5.1. Let us look at how to formalize a variation of Example 1.2. The specification $\Sigma(B, V)$ includes two conjuncts $C_1(B, V)$ and $C_2(B, V)$. $C_1(B, V)$ states that every pair $\langle k, e \rangle$ of V corresponds to a $\langle k, S \rangle$ in B with $e \in S$:

$$\forall v \in V \; \exists b \in B. \; \pi_1(v) =_{\mathfrak{U}} \pi_1(b) \wedge \pi_2(v) \; \hat{\in} \; \pi_2(b)$$

$C_2(B, V)$ is:

$$\forall b \in B \; \forall e \in \pi_2(b) \; \exists v \in V. \; \pi_1(v) =_{\mathfrak{U}} \pi_1(b) \wedge \pi_2(v) =_{\mathfrak{U}} e$$

Let us assume a stronger constraint, $\Sigma_{\mathsf{lossless}}(B)$, saying that the first component is a key and second is non-empty:

$$\forall b \in B \; \forall b' \in B. \pi_1(b) =_{\mathfrak{U}} \pi_1(b') \to b \equiv b'$$
$$\wedge \; \forall b \in B \; \exists e \in \pi_2(b). \top$$

With $\Sigma_{\mathsf{lossless}}$ we can show something stronger than in Example 1.2: $\Sigma \wedge \Sigma_{\mathsf{lossless}}$ implicitly defines B in terms of V. That is, the view determines the identity query, which is witnessed by a proof of

$$\Sigma(B, V) \wedge \Sigma_{\mathsf{lossless}}(B) \wedge \Sigma(B', V) \wedge \Sigma_{\mathsf{lossless}}(B') \to B \equiv B'$$

Let's prove this informally. Assuming the premise, it is sufficient to prove $B \subseteq B'$ by symmetry. So fix $\langle k, S \rangle \in B$. By the second conjunct of $\Sigma_{\mathsf{lossless}}(B)$, we know there is $e \in S$. Thus by $C_2(B, V)$, V contains the pair $\langle k, e \rangle$. Then, by $C_1(B', V)$, there is a S' such that $\langle k, S' \rangle \in B'$. To conclude it suffices to show that $S \equiv S'$. There are two similar directions, let us detail the inclusion $S \subseteq S'$; so fix $s \in S$. By $C_2(B, V)$, we have $\langle k, s \rangle \in V$. By $C_1(B', V)$ there exists S'' such that $\langle k, S'' \rangle \in B'$ with $s \; \hat{\in} \; S''$. But since we also have $\langle k, S' \rangle \in B'$, the constraint $\Sigma_{\mathsf{lossless}}(B)$ implies that $S' \equiv S''$, so $s \in S'$ as desired.

Effective Beth result. A derivation of implicit definability in our proof system will be referred to as a *witness* to the implicit definability of o in terms of $\vec{i}$ relative to φ. Formally, this is a derivation witnessing the judgement:

$$\varphi(\vec{i}, \vec{a}, o) \; \wedge \; \varphi(\vec{i}, \vec{a}', o') \vdash o \equiv_T o'$$

With these definitions, we now state formally our main result, the effective version of Corollary 4.3.

Theorem 5.2 (Effective implicit to explicit for nested data)**.** *Given a witness to the implicit definition of o in terms of $\vec{i}$ relative to Δ_0 $\varphi(\vec{i}, \vec{a}, o)$, one can compute NRC expression E such that for any $\vec{i}$, $\vec{a}$ and o, if $\varphi(\vec{i}, \vec{a}, o)$ then $E(\vec{i}) = o$. Furthermore, if the witness is EL-normalized, this can be done in polynomial time.*

Application to views and queries. We now state the consequence for effective rewriting queries over views mentioned in the introduction. Consider a query given by NRC expression E_Q over inputs $\vec{B}$ and NRC expressions $E_{V_1} \ldots E_{V_n}$ over $\vec{B}$. E_Q is *determined* by $E_{V_1} \ldots E_{V_n}$, if every two nested relations (finite or infinite) interpreting $\vec{B}$ that agree on the output of each E_{V_i} agree on the output of E_Q. An NRC rewriting of E_Q in terms of $E_{V_1} \ldots E_{V_n}$ is an expression $R(V_1 \ldots V_n)$ such that for any nested relation $\vec{B}$, if we evaluate each E_{V_i} on $\vec{B}$ to obtain V_i and evaluate R on the resulting $V_1 \ldots V_n$, we obtain $Q(\vec{B})$.

Given E_Q and $E_{V_1} \ldots E_{V_n}$, let $\Sigma_{\vec{V},Q}(\vec{V}, \vec{B}, Q, \ldots)$ conjoin the input-output specifications, as defined in Section 3, for $E_{V_1} \ldots E_{V_n}$ and E_Q. This formula has variables $\vec{B}, V_1 \ldots V_n, Q$ along with auxiliary variables for subqueries. A proof witnessing determinacy of E_Q by $E_{V_1} \ldots E_{V_n}$, is a proof that $\Sigma_{V,Q}$ implicitly defines Q in terms of $\vec{V}$.

Corollary 5.3. *From a witness that a set of* NRC *views* $\vec{V}$ *determines an* NRC *query* Q, *we can produce an* NRC *rewriting of* Q *in terms of* $\vec{V}$. *If the witness is* EL-*normalized, this can be done in* PTIME.

The notion of determinacy of a query over views relative to a Δ_0 theory (e.g. the key constraint in Example 1.2) is a straightforward generalization of the definitions above, and Corollary 5.3 extends to this setting.

In the case where we are dealing with flat relations, the effective version is well-known: see Toman and Weddell's [TW11], and the discussion in [FKN13, BCLT16].

We emphasize that the result involves equivalence up to extensionality, which underlines the distinction from the classical Beth theorem. If we wrote out implicit definability up to extensionality as an entailment involving two copies of the signature, we would run into problems in applying the standard proof of Beth's theorem.

5.2. **Tools for the effective Beth theorem.**

5.2.1. *Interpolation.* The first tool for our effective Beth theorem, Theorem 5.2, will be an effective version of interpolation Proposition 4.4. Recall that interpolation results state that if we have an entailment involving two formulas, a "left" formula φ_L and a "right" formula φ_R, we can get an "explanation" for the entailment that factors through an expression only involving non-logical symbols (in our case, variables) that are common to φ_L and φ_R.

Theorem 5.4. *Let* Θ *be an* $\in$-*context and* Γ, Δ *finite multisets of* Δ_0 *formulas. Then from any proof of* $\Theta;\ \Gamma \vdash \Delta$, *we can compute in linear time an extended* Δ_0 *formula* θ *with* $FV(\theta) \subseteq FV(\Theta, \Gamma) \cap FV(\Delta)$ *such that* $\Theta;\ \Gamma \vdash \theta$ *and* $\emptyset;\ \theta \vdash \Delta$.

The θ produced by the theorem is a *Craig interpolant*. Craig's interpolation theorem [Cra57a] states that when $\Gamma \vdash \Delta$ with Γ, Δ in first-order logic, such a θ exists in first-order logic. Our variant states one can find θ in Δ_0 efficiently from a proof of the entailment in either of our Δ_0 proof systems. We have stated the result for the 2-sided system. It holds also for the EL-normalized system, where the partition of the formulas into left and right of the proof symbol is arbitrary. The argument is induction on proof length, roughly following prior interpolation algorithms [Smu68b]. See Appendix H.

We compare with the model-theoretic statement Proposition 4.4. There, the interpolant is Δ_0, while here in the effective version it is *extended* Δ_0. This is due to the linear time

requirement, which leads to the involvement of equalities $=_T$ in the interpolation algorithm. The construction would also work in plain Δ_0 if $\equiv_T$ is used instead. However, $\equiv_T$ is a shorthand for a formula whose size is exponential in the term depth of the type T.

5.2.2. *Some admissible rules.* As we mentioned earlier, our EL-normalized proof system is extremely low-level, and so it is convenient to have higher-level proof rules as macros. We formalize this below.

Definition 5.5. A rule with premise $\Theta' \vdash \Delta'$ and conclusion $\Theta \vdash \Delta$

$$\frac{\Theta' \vdash \Delta'}{\Theta \vdash \Delta}$$

is *(polytime) admissible* in a given calculus if a proof of the conclusion $\Theta \vdash \Delta$ in that calculus can be computed from a proof of the premise $\Theta' \vdash \Delta'$ (in polynomial time).

Up to rewriting the sequent to be one-sided, all the rules in Figure 2 are polytime admissible in the EL-normalized calculus. Our main theorem will rely on the polytime admissibility within the EL-normalized calculus of additional rules that involve chains of existential quantifiers. To state them, we need to introduce a generalization of bounded quantification: "quantifying over subobjects of a variable".

Definition 5.6. *For every type T, define a subset of the non-empty words over the three-letter alphabet $\{1, 2, \ni\}$ of* subtype occurrences *of T inductively as follows:*
- *If p is a subtype occurrence of T or the empty word then the concatenation $\ni, p$ is a subtype occurrence of $\mathsf{Set}(T)$.*
- *If $i \in \{1, 2\}$ and p is a subtype occurrence of T_i, i, p is a subtype occurrence of $T_1 \times T_2$.*

Given subtype occurrence p and quantifier symbol $\mathbf{Q} \in \{\forall, \exists\}$, define the notation $\mathbf{Q}\, x \in_p t.\varphi$ by induction on p:
- *$\mathbf{Q}\, x \in_\ni t.\varphi$ is $\mathbf{Q}\, x \in t$*
- *$\mathbf{Q}\, x \in_{\ni,p} t.\varphi$ is $\mathbf{Q}\, y \in t.\mathbf{Q}\, x \in_p y.\varphi$ with y a fresh variable*
- *$\mathbf{Q}\, x \in_{i,p} t.\varphi$ is $\mathbf{Q}\, x \in_p \pi_i(t).\varphi$ when $i \in \{1, 2\}$.*

Now we are ready to state the results we need on admissibility, referring in each case to the EL-normalized calculus. Some further rules, that are easily seen to be admissible, are used in the appendices. All proofs are found in Appendix I. The first states that if we have proven that there exists a subobject of o' equivalent to object r, then we can prove that for each element z of r there is a corresponding equivalent subobject z' within o'. Furthermore, this can be done effectively, and the output proof has at most the same size: that is, there is not even a polynomial blow-up involved.

Lemma 5.7. *Assume p is a subtype occurrence for the type of the term o'. The following is polytime admissible*

$$\frac{\Theta \vdash \Delta,\ \exists r' \in_p o'.\ r \equiv_{\mathsf{Set}(T')} r'}{\Theta, z \in r \vdash \Delta,\ \exists z' \in_{\ni,p} o'.\ z \equiv_{T'} z'}$$

Furthermore, the size of the output proof is at most the size of the input proof.

The second result states that we can move between an equivalence of r, r' and a universally-quantified biconditional between memberships in r and r'. Because we are dealing with Δ_0 formulas, the universal quantification has to be bounded by some additional variable a.

Lemma 5.8. *The following is polytime admissible (where p is a subtype occurrence of the type of o')*

$$\frac{\Theta \vdash \Delta, \exists r' \in_p o'.\ r \equiv_{\mathsf{Set}(T')} r'}{\Theta \vdash \Delta, \exists r' \in_p o'.\forall z \in a.\ z \hat{\in} r \leftrightarrow z \hat{\in} r'}$$

5.2.3. *The* NRC *Parameter Collection Theorem.* Our last tool is a kind of interpolation result connecting Δ_0 formulas and NRC:

Theorem 5.9 (NRC Parameter Collection). *Let L, R be sets of variables with $C = L \cap R$ and*

- φ_L *and* $\lambda(z)$ Δ_0 *formulas over* L
- φ_R *and* $\rho(z, y)$ Δ_0 *formulas over* R
- r *a variable of* R *and* c *a variable of* C.

    Suppose that we have an EL-normalized proof of

$$\varphi_L \wedge \varphi_R \ \rightarrow \ \exists y \in_p r\ \forall z \in c\ (\lambda(z) \leftrightarrow \rho(z, y))$$

    Then one may compute in polynomial time an NRC *expression E with free variables in C such that*

$$\varphi_L \wedge \varphi_R \rightarrow \{z \in c \mid \lambda(z)\} \in E$$

If λ was a "common formula" — one using only variables in C — then the nested relation $\{z \in c \mid \lambda(z)\}$ would be definable over C in NRC via Δ_0-comprehension. Unfortunately λ is a "left formula", possibly with variables outside of C. Our hypothesis is that it is equivalent to a "parameterized right formula": a formula with variables in R and parameters that lie below them. Intuitively, this can happen only if λ can be rewritten to a formula $\rho'(z, x)$ with variables of C and a distinguished $c_0 \in C$ such that

$$\varphi_L \wedge \varphi_R \ \rightarrow \ \exists x \in_p c_0\ \forall z \in c\ (\lambda(z) \leftrightarrow \rho'(z, x))$$

And if this is true, we can use an NRC expression over C to define a set that will contain the correct "parameter" value x defining λ. From this we can define a set containing the nested relation $\{z \in c \mid \lambda(z)\}$. A formalization of this rough intuition — "when left formulas are equivalent to parameterized right formulas, they are equivalent to parameterized common formulas" — is given in Section K, where a similar statement that does not mention NRC is proven for first-order logic (sadly it does not seem strong enough to derive Theorem 5.9).

    We now give the proof of the theorem.

    To get the desired conclusion, we need to prove a more general statement by induction over proof trees. Besides making the obvious generalization to handle two multisets of formulas instead of the particular formulas φ_L and φ_R, as well as some corresponding left and right $\in$-contexts, that may appear during the proof, we need to additionally generate a new formula θ that only uses common variables, which can replace φ_R in the conclusion. This is captured in the following lemma:

Lemma 5.10. *Let L, R be sets of variables with $C = L \cap R$ and*

- $\Delta_L, \lambda(z)$ *a multiset of Δ_0 formulas over L*
- $\Delta_R, \rho(z, y)$ *a multiset of Δ_0 formulas over R*
- Θ_L *(respectively Θ_R) an $\in$-context over L (respectively over R)*
- $r_1, \ldots, r_k$ *variables of R and c a variable of C.*

- *we write $\mathcal{G}(r_i)$ for $\exists y \in_{p_i} r_i \; \forall z \in c \; (\lambda(z) \leftrightarrow \rho(z,y))$*

  Suppose that we have a EL-normalized proof tree with conclusion

  $$\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \mathcal{G}(r_1), \ldots, \mathcal{G}(r_k)$$

  Then one may compute in polynomial time an NRC expression E and an extended Δ_0 formula θ using only variables from C such that

  $$\Theta_L \models_{\mathsf{nested}} \Delta_L, \theta \vee \{z \in c \mid \lambda(z)\} \in E \quad and \quad \Theta_R \models_{\mathsf{nested}} \Delta_R, \neg\theta$$

The theorem follows easily from this lemma, so we focus on proving the lemma, by induction over the size of the proof of $\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \mathcal{G}(r_1), \ldots, \mathcal{G}(r_k)$, making a case distinction according to which rule is applied last. The way θ will be built will, perhaps unsurprisingly, be very reminiscent of the way interpolants are normally constructed in standard proof systems [Fit96, Smu68b].

For readability, we adopt the following conventions:

- We write $\tilde{\mathcal{G}}$ for the multiset of formulas $\mathcal{G}(r_1), \ldots, \mathcal{G}(r_k)$ and Λ for the expression $\{z \in c \mid \lambda(z)\}$.
- For formulas and NRC expressions obtained by applying the induction hypothesis, we use the names θ^{IH} and E^{IH} (or $\theta_1^{\mathsf{IH}}, \theta_2^{\mathsf{IH}}$ and $E_1^{\mathsf{IH}}, E_2^{\mathsf{IH}}$ when the induction hypothesis is applied several times). In each subcase, our goal will be to build suitable θ and E.
- We will color pairs of terms, formulas and multisets of formulas according to whether they are part of either $\Theta_L; \Delta_L$ or $\Theta_R; \Delta_R$ either at the start of the case analysis or when we want to apply the induction hypothesis. In particular, the last sequent of the proof under consideration will be depicted as

  $$\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \mathcal{G}$$

- Unless it is non-trivial, we leave checking that the free variables in our proposed definition for E and Θ are taken among variables of C to the reader.

In two of the cases below, we will make use of some syntactic sugar on top of bounded quantification. We introduce

$$\exists x_1 \ldots x_n | t \in b . \varphi \quad \text{and} \quad \forall x_1 \ldots x_n | t \in b . \varphi$$

as notation. This will be an extended Δ_0 formula, intuitively quantifying over variables $x_1, \ldots, x_n$ that occur in tuple-term t, which is bounded by set b. Here φ is an extended Δ_0 formula. The variables other than $x_1, \ldots, x_n$ occurring in t remain free in the resulting formula.

To make this precise, let $\mathcal{R}ef(t, v, u)$ be the set of the terms expressed with projection applied to u that refer to an occurrence of variable v in tuple-term t, ordered according to the occurrences when t is printed. Formally, $\mathcal{R}ef(t, t, u) := \{u\}$; $\mathcal{R}ef(t, v, u) := \{\}$ if t is a variable other than v; and $\mathcal{R}ef(\langle t_1, t_2 \rangle, v, u) := \mathcal{R}ef(t_1, v, \pi_1(u)) \cup \mathcal{R}ef(t_2, v, \pi_2(u))$. We write $\mathcal{R}ef_i(t, v, u)$ for the i-th member of $\mathcal{R}ef(t, v, u)$. We can then define $\exists x_1 \ldots x_n | t \in b . \varphi$ as

$$\exists y \in b . \varphi[\mathcal{R}ef_1(t, x_1, y)/x_1, \ldots, \mathcal{R}ef_1(t, x_n, y)/x_n] \; \wedge$$
$$\bigwedge_{z \text{ is a variable occurring in } t \text{ other than } x_1, \ldots, x_n} z = \mathcal{R}ef_1(t, z, y) \; \wedge$$
$$\bigwedge_{u \text{ is a variable occurring in } t \text{ and } \mathcal{R}ef_j(t, y, u) \text{ for } j > 1 \text{ in } \mathcal{R}ef(t, y, u)} \mathcal{R}ef_1(t, u, y) = \mathcal{R}ef_j(t, u, y)$$

$\forall x_1 \ldots x_n | t \in b . \varphi$ can be defined analogously.

Here equality $=_T$ (written above without the type decoration) of *extended* Δ_0 formulas comes into play, in contrast to $\equiv_T$, to meet the polynomial time requirements of Theorems 5.4 and 5.9 as indicated on p. 27.

As an example for the notation consider

$$\exists x_1 x_2 | \langle \langle x_1, x_2 \rangle, z \rangle \in b . x_1 \equiv x_2,$$

which stands for

$$\exists y \in b . \pi_1(\pi_1(y)) \equiv \pi_2(\pi_1(y)) \wedge z = \pi_2(y).$$

With these conventions in mind, let us proceed.

- If the last rule applied is the $\top$ rule, in both cases we are going to take $E := \emptyset$, but pick θ to be $\bot$ or $\top$ according to whether $\top$ occurs in Δ_L or Δ_R; we leave checking the details to the reader.
- If the last rule applied is the $\wedge$ rule, we have two cases according to the position of the principal formula $\varphi_1 \wedge \varphi_2$. In both cases, E will be obtained by unioning NRC expressions obtained from the induction hypothesis, and θ will be either a disjunction or a conjunction.
  - If we have $\Delta_L = \varphi_1 \wedge \varphi_2, \Delta'_L$, so that the proof has shape

$$\frac{\Theta_L, \Theta_R \vdash \varphi_1, \Delta'_L, \Delta_R, \tilde{\mathcal{G}} \qquad \Theta_L, \Theta_R \vdash \varphi_2, \Delta'_L, \Delta_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \varphi_1 \wedge \varphi_2, \Delta'_L, \Delta_R, \tilde{\mathcal{G}}}$$

by the induction hypothesis, we have NRC expressions $E_1^{\mathsf{IH}}, E_2^{\mathsf{IH}}$ and formulas $\theta_1^{\mathsf{IH}}, \theta_2^{\mathsf{IH}}$ such that

$$\begin{array}{lll} \Theta_L \models_{\mathsf{nested}} \varphi_1, \Delta'_L, \theta_1^{\mathsf{IH}} \vee \Lambda \in E_1^{\mathsf{IH}} & \text{and} & \Theta_R \models_{\mathsf{nested}} \Delta_R, \neg\theta_1^{\mathsf{IH}} \\ \Theta_L \models_{\mathsf{nested}} \varphi_2, \Delta'_L, \theta_2^{\mathsf{IH}} \vee \Lambda \in E_1^{\mathsf{IH}} & \text{and} & \Theta_R \models_{\mathsf{nested}} \Delta_R, \neg\theta_2^{\mathsf{IH}} \end{array}$$

In that case, we take $E = E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}}$ and $\theta := \theta_1^{\mathsf{IH}} \vee \theta_2^{\mathsf{IH}}$. Weakening the properties on the left column, we have

$$\Theta_L \models_{\mathsf{nested}} \varphi_i, \Delta'_L, \theta \vee \Lambda \in E$$

for both $i \in \{1, 2\}$, so we have

$$\Theta_L \models_{\mathsf{nested}} \varphi_1 \wedge \varphi_2, \Delta'_L, \theta \vee \Lambda \in E$$

as desired. Since $\neg\theta = \neg\theta_1^{\mathsf{IH}} \wedge \neg\theta_2^{\mathsf{IH}}$, we get

$$\Theta_R \models_{\mathsf{nested}} \Delta_R, \neg\theta$$

by combining both properties from the right column.
  - The dual case where $\Delta_R = \varphi_1 \wedge \varphi_2, \Delta'_R$ is handled similarly, except that we set $\theta := \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}}$.
- Suppose the last rule applied is $\vee$ with principal formula $\varphi_1 \vee \varphi_2$. Depending on whether $\Delta_L = \varphi_1 \vee \varphi_2, \Delta'_L$ or $\Delta_R = \varphi_1 \vee \varphi_2, \Delta'_R$, the proof will end with one of the following steps

$$\frac{\Theta_L, \Theta_R \vdash \varphi_1, \varphi_2, \Delta'_L, \Delta_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \varphi_1 \vee \varphi_2, \Delta'_L, \Delta_R, \tilde{\mathcal{G}}} \qquad \text{or} \qquad \frac{\Theta_L, \Theta_R \vdash \Delta'_L, \varphi_1, \varphi_2, \Delta_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \Delta'_L, \varphi_1 \vee \varphi_2, \Delta_R, \tilde{\mathcal{G}}}$$

In both cases, we apply the inductive hypothesis according to the obvious splitting of contexts and multisets of formulas, to get an NRC definition E^{IH} along with a formula θ^{IH} that satisfy the desired semantic property. We set $E := E^{\mathsf{IH}}$ and $\theta := \theta^{\mathsf{IH}}$.

- Suppose the last rule applied is $\forall$ with principal formula $\forall x \in b.\varphi$. As in the previous case, depending on whether $\Delta_L = \forall x \in b.\varphi, \Delta'_L$ or $\Delta_R = \forall x \in b.\varphi, \Delta'_R$, the proof will end with one of the following steps (assuming y is fresh below)

$$\frac{\Theta_L, y \in b, \Theta_R \vdash \varphi[y/x], \Delta'_L, \Delta_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \forall x \in b.\varphi, \Delta'_L, \Delta_R, \tilde{\mathcal{G}}} \quad \text{or} \quad \frac{\Theta_L, \Theta_R, y \in b \vdash \Delta'_L, \varphi[y/x], \Delta_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \Delta'_L, \forall x \in b.\varphi, \Delta_R, \tilde{\mathcal{G}}}$$

  In both cases, we again apply the inductive hypothesis according to the obvious splitting of contexts and multisets of formulas to get an NRC definition E^{IH} and a formula θ^{IH} that satisfy the desired semantic property. We set $E := E^{\mathsf{IH}}$ and $\theta := \theta^{\mathsf{IH}}$.

- Now we consider the case where the last rule applied is $\exists$. Here we have two main subcases, according to whether the formula that is instantiated in the premise, that is, the principal formula, belongs to $\tilde{\mathcal{G}}$ or not. In the first case, we have two further subcases according to whether the instantiated formula still has a leading existential quantifier or not.

  - If the principal formula is of the shape $\mathcal{G} = \exists y \in r \; \forall z \in c. \; (\lambda(z) \leftrightarrow \rho(z,y))$ (so in particular, $\mathcal{G}$ has a single leading existential quantifier) and is a member of $\tilde{\mathcal{G}}$, the proof necessarily has shape

$$\underset{\wedge}{\overset{\vee}{\frac{\dfrac{\Theta_L, \Theta_R, x \in c \vdash \Delta_L, \Delta_R, \neg\rho(x,w), \lambda(x), \tilde{\mathcal{G}}}{\Theta_L, \Theta_R, x \in c \vdash \Delta_L, \Delta_R, \rho(x,w) \to \lambda(x), \tilde{\mathcal{G}}} \quad \overset{\vee}{\dfrac{\Theta_L, \Theta_R, x \in c \vdash \Delta_L, \Delta_R, \neg\lambda(x), \rho(x,w), \tilde{\mathcal{G}}}{\Theta_L, \Theta_R, x \in c \vdash \Delta_L, \Delta_R, \lambda(x) \to \rho(x,w), \tilde{\mathcal{G}}}}}{\underset{\exists}{\dfrac{\underset{\forall}{\dfrac{\Theta_L, \Theta_R, x \in c \vdash \Delta_L, \Delta_R, \lambda(x) \leftrightarrow \rho(x,w), \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \forall z \in c. \; (\lambda(z) \leftrightarrow \rho(z,w)), \tilde{\mathcal{G}}}}}{\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \tilde{\mathcal{G}}}}}}$$

  where x is a fresh variable So in particular, we have two strict subproofs with respective conclusions

$$\begin{aligned}
&\Theta_L, x \in c, \Theta_R \; \vdash \; \lambda(x), \Delta_L, \neg\rho(x,w), \Delta_R, \tilde{\mathcal{G}} \\
\text{and} \quad &\Theta_L, x \in c, \Theta_R \; \vdash \; \neg\lambda(x), \Delta_L, \rho(x,w), \Delta_R, \tilde{\mathcal{G}}
\end{aligned}$$

  Applying the inductive hypothesis, we obtain NRC expressions E_1^{IH}, E_2^{IH} and formulas $\theta_1^{\mathsf{IH}}, \theta_2^{\mathsf{IH}}$ which contain free variables in $C \cup \{x\}$ such that all of the following hold

$$\Theta_L, x \in c \models_{\mathsf{nested}} \quad \lambda(x), \Delta_L, \theta_1^{\mathsf{IH}} \vee \Lambda \in E_1^{\mathsf{IH}} \tag{5.2}$$

$$\text{and} \quad \Theta_L, x \in c \models_{\mathsf{nested}} \neg\lambda(x), \Delta_L, \theta_2^{\mathsf{IH}} \vee \Lambda \in E_2^{\mathsf{IH}} \tag{5.3}$$

$$\text{and} \quad \Theta_R \models_{\mathsf{nested}} \neg\rho(x,w), \Delta_R, \neg\theta_1^{\mathsf{IH}} \tag{5.4}$$

$$\text{and} \quad \Theta_R \models_{\mathsf{nested}} \quad \rho(x,w), \Delta_R, \neg\theta_2^{\mathsf{IH}} \tag{5.5}$$

  With this in hand, we set

$$\theta := \exists x \in c. \; \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}} \qquad \text{and} \qquad E := \left\{\left\{ x \in c \mid \theta_2^{\mathsf{IH}} \right\}\right\} \; \cup \; \bigcup \left\{ E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}} \mid x \in c \right\}$$

  Note in particular that the free variables of E and θ are contained in C, since we bind x. The bindings of x have radically different meaning across the two main components $E_1 := \{\{x \in c \mid \theta_2^{\mathsf{IH}}\}\}$ and $E_2 := \{E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}} \mid x \in c\}$ of $E = E_1 \cup E_2$. E_1 consists of a single definition corresponding to the restriction of c to θ_2^{IH}, and there x plays the role of an element being defined. On the other hand, E_2 corresponds to the joining of all the definitions obtained inductively, which may contain an $x \in c$ as a parameter. So we have two families of potential definitions for Λ indexed by $x \in c$ that we join together.

Now let us show that we have the desired semantic properties. First we need to show that E contains a definition for Λ under the right hypotheses, i.e.,

$$\Theta_L \models_{\mathsf{nested}} \Delta_L, \exists x \in c.\ \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}}, \Lambda \in \left(\{\{x \in c \mid \theta_2^{\mathsf{IH}}\}\} \cup \bigcup \{E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}} \mid x \in c\} \right) \qquad (5.6)$$

which can be rephrased as

$$\Theta_L \models_{\mathsf{nested}} \Delta_L, \exists x \in c.\ \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}}, \Lambda = \left\{ x \in c \mid \theta_2^{\mathsf{IH}} \right\}, \exists x \in c.\ \Lambda \in E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}}$$

Now concentrate on the statement $\Lambda = \left\{ x \in c \mid \theta_2^{\mathsf{IH}} \right\}$. It would follow from the two inclusions $\Lambda \subseteq \left\{ x \in c \mid \theta_2^{\mathsf{IH}} \right\}$ and $\left\{ x \in c \mid \theta_2^{\mathsf{IH}} \right\} \subseteq \Lambda$, so, recalling that $\Lambda = \{x \in c \mid \lambda(x)\}$, the overall conclusion would follow from having

$$\Theta_L, x \in c \models_{\mathsf{nested}} \Delta_L, \exists x \in c.\ \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}}, \lambda(x) \to \theta_2^{\mathsf{IH}}, \exists x \in c.\ \Lambda \in E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}}$$

and $\qquad \Theta_L, x \in c \models_{\mathsf{nested}} \Delta_L, \exists x \in c.\ \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}}, \theta_2^{\mathsf{IH}} \to \lambda(x), \exists x \in c.\ \Lambda \in E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}}$

Those in turn follow from the following two statements

$$\Theta_L, x \in c \models_{\mathsf{nested}} \Delta_L, \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}}, \neg\lambda(x), \theta_2^{\mathsf{IH}}, \Lambda \in E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}}$$

and $\qquad \Theta_L, x \in c \models_{\mathsf{nested}} \Delta_L, \theta_1^{\mathsf{IH}} \wedge \theta_2^{\mathsf{IH}}, \neg\theta_2^{\mathsf{IH}}, \lambda(x), \Lambda \in E_1^{\mathsf{IH}} \cup E_2^{\mathsf{IH}}$

which are straightforward consequences of 5.3 and 5.2 respectively. This concludes the proof of 5.6.

Now we only need to prove a final property, which is

$$\Theta_R \models_{\mathsf{nested}} \Delta_R, \forall x \in c.\ \neg\theta_1^{\mathsf{IH}} \vee \neg\theta_2^{\mathsf{IH}}$$

which is equivalent to the validity of

$$\Theta_R, x \in c \models_{\mathsf{nested}} \Delta_R, \neg\theta_1^{\mathsf{IH}} \vee \neg\theta_2^{\mathsf{IH}}$$

which can be obtained by combining 5.4 and 5.5 with excluded middle for $\rho(x, w)$.

– If the principal formula is of the shape $\mathcal{G} = \exists r' \in r\ \mathcal{G}'$ where $\mathcal{G}'$ begins with another existential quantifier and $\mathcal{G}$ is a member of $\tilde{\mathcal{G}} = \mathcal{G}, \tilde{\mathcal{G}}'$, the proof necessarily has shape

$$\frac{\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \mathcal{G}', \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \tilde{\mathcal{G}}}$$

then we can apply the induction hypothesis where in lieu of $\tilde{\mathcal{G}}$ we have $\mathcal{G}', \tilde{\mathcal{G}}$ and obtain θ^{IH} and E^{IH}. It is then clear that we can simply set $\theta := \theta^{\mathsf{IH}}$ and $E := E^{\mathsf{IH}}$.

– If the principal formula is not a member of $\tilde{\mathcal{G}}$, then we have two subcases corresponding to whether the principal formula under consideration occurs in Δ_L or Δ_R, and whether the relevant membership statement that witnesses the instantiation is a member of Θ_L or Θ_R. Let us list all of the different alternatives

∗ If the last step of the proof has shape

$$\frac{\Theta_L', w \in b, \Theta_R \vdash \Delta_L', \varphi[w/x], \exists x \in b\ \varphi, \Delta_R, \tilde{\mathcal{G}}}{\Theta_L', w \in b, \Theta_R \vdash \Delta_L', \exists x \in b\ \varphi, \Delta_R, \tilde{\mathcal{G}}}$$

with $\Theta_L = \Theta_L', w \in b$ and $\Delta_L = \Delta_L', \exists x \in b\ \varphi$, we can conclude by setting $\theta := \theta^{\mathsf{IH}}$ and $E := E^{\mathsf{IH}}$, essentially because the set of free variables L, R and C can be taken to be the same in the premise.

∗ In the dual case where the last step has shape

$$\frac{\Theta_L, \Theta'_R, w \in b \vdash \Delta_L, \Delta'_R, \varphi[w/x], \exists x \in b\ \varphi, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R, w \in b \vdash \Delta_L, \Delta'_R, \exists x \in b\ \varphi, \tilde{\mathcal{G}}}$$

with $\Theta_R = \Theta'_R, w \in b$ and $\Delta_R = \Delta'_R, \exists x \in b\ \varphi$, we can also conclude immediately by setting $\theta := \theta^{\mathsf{IH}}$ and $E := E^{\mathsf{IH}}$.

∗ If the last step has shape

$$\frac{\Theta'_L, w \in b, \Theta_R \vdash \Delta_L, \Delta'_R, \varphi[w/x], \exists x \in b\ \varphi, \tilde{\mathcal{G}}}{\Theta'_L, w \in b, \Theta_R \vdash \Delta_L, \Delta'_R, \exists x \in b\ \varphi, \tilde{\mathcal{G}}}$$

with $\Theta_L = \Theta'_L, w \in b$ and $\Delta_R = \Delta'_R, \exists x \in b\ \varphi$, we need to do something non-trivial. We can still use the inductive hypothesis to obtain θ^{IH} and E^{IH}, but they may feature variables $x_1, \ldots, x_n$ of w that are in L as free variables. But we also have that the free variables of b are included in C. With that in mind, we can set $\theta := \exists x_1 \ldots x_n | w \in b\, . \, \theta^{\mathsf{IH}}$ and $E := \bigcup \{E^{\mathsf{IH}} \mid w \in b\}$.

∗ If the last step has shape

$$\frac{\Theta_L, \Theta_R, w \in b \vdash \Delta'_L, \varphi[w/x], \exists x \in b\ \varphi, \Delta_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta'_R, w \in b \vdash \Delta'_L, \exists x \in b\ \varphi, \Delta_R, \tilde{\mathcal{G}}}$$

with $\Theta_R = \Theta'_R, w \in b$ and $\Delta_L = \Delta'_L, \exists x \in b\ \varphi$, we proceed similarly by setting $\theta := \forall x_1 \ldots x_n | w \in b\, . \, \theta^{\mathsf{IH}}$ and $E := \bigcup \{E^{\mathsf{IH}} \mid w \in b\}$.

- The case of the $=$ rule can be handled exactly as the $\top$ rule.
- For the $\neq$ rule, we distinguish several subcases:
  - If we have $\Delta_L = y \neq_{\mathfrak{U}} z, \alpha[y/x], \Delta'_L$ or $\Delta_R = y \neq_{\mathfrak{U}} z, \alpha[y/x], \Delta'_R$, so that the last step has one of the two following shapes

$$\frac{\Theta_L, \Theta_R \vdash y \neq_{\mathfrak{U}} z, \alpha[y/x], \alpha[z/x], \Delta'_L, \Delta'_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash y \neq_{\mathfrak{U}} z, \alpha[y/x], \Delta'_L, \Delta'_R, \tilde{\mathcal{G}}} \qquad \frac{\Theta_L, \Theta_R \vdash \Delta'_L, y \neq_{\mathfrak{U}} z, \alpha[y/x], \alpha[z/x], \Delta'_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \Delta'_L, y \neq_{\mathfrak{U}} z, \alpha[y/x], \Delta'_R, \tilde{\mathcal{G}}}$$

we can apply the induction hypothesis to obtain some θ^{IH} and E^{IH} such that setting $\theta := \theta^{\mathsf{IH}}$ and $E := E^{\mathsf{IH}}$ solves this subcase; we leave checking the additional properties to the reader.

  - Otherwise, if we have $\Delta_L = y \neq_{\mathfrak{U}} z, \Delta'_L$, $\Delta_R = \alpha[y/x], \Delta'_R$ and a last step of shape

$$\frac{\Theta_L, \Theta_R \vdash y \neq_{\mathfrak{U}} z, \Delta'_L, \alpha[y/x], \alpha[z/x], \Delta'_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash y \neq_{\mathfrak{U}} z, \Delta'_L, \alpha[y/x], \Delta'_R, \tilde{\mathcal{G}}}$$

In that case, the inductive hypothesis gives θ^{IH} and E^{IH} with free variables in $C \cup \{z\}$ such that

$$\Theta_L \models_{\mathsf{nested}} y \neq_{\mathfrak{U}} z, \Delta'_L, \theta^{\mathsf{IH}}, \Lambda \in E^{\mathsf{IH}} \quad \text{and} \quad \Theta_R \models_{\mathsf{nested}} \alpha[y/x], \alpha[z/x], \Delta'_R, \neg\theta^{\mathsf{IH}}$$

We then have two subcases according to whether $z \in C$ or not

∗ If $z \in C$, we can take $\theta := \theta^{\mathsf{IH}} \wedge y =_{\mathfrak{U}} z$ and $E := E^{\mathsf{IH}}$. Their free variables are in C and we only need to check

$$\Theta_L \models_{\mathsf{nested}} y \neq_{\mathfrak{U}} z, \Delta'_L, \theta^{\mathsf{IH}} \wedge y =_{\mathfrak{U}} z, \Lambda \in E^{\mathsf{IH}} \quad \text{and} \quad \Theta_R \models_{\mathsf{nested}} \alpha[y/x], \Delta'_R, \neg\theta^{\mathsf{IH}}, y \neq_{\mathfrak{U}} z$$

which follow easily from the induction hypothesis.

* Otherwise, we take $\theta := \theta^{\mathsf{IH}}[y/z]$ and $E := E^{\mathsf{IH}}[y/z]$. In that case, note that we have $\alpha[z/x][y/z] = \alpha[y/x]$ (which would not be necessarily the case if z belonged to C). This allows to conclude that we have

$$\Theta_L \models_{\mathsf{nested}} y \neq_{\mathfrak{U}} z, \Delta'_L, \theta^{\mathsf{IH}}[y/z], \Lambda \in E^{\mathsf{IH}}[y/z] \quad \text{and} \quad \Theta_R \models_{\mathsf{nested}} \alpha[y/x], \Delta'_R, \neg\theta^{\mathsf{IH}}[y/z]$$

directly from the induction hypothesis.

– Otherwise, if we have $\Delta_L = \alpha[y/x], \Delta'_L$, $\Delta_R = y \neq_{\mathfrak{U}} z, \Delta'_R$ and a last step of shape

$$\frac{\Theta_L, \Theta_R \vdash \alpha[y/x], \alpha[z/x], \Delta'_L, y \neq_{\mathfrak{U}} z, \Delta'_R, \tilde{\mathcal{G}}}{\Theta_L, \Theta_R \vdash \alpha[y/x], \Delta'_L, y \neq_{\mathfrak{U}} z, \Delta'_R, \tilde{\mathcal{G}}}$$

In that case, the inductive hypothesis gives θ^{IH} and E^{IH} with free variables in $C \cup \{z\}$ such that

$$\Theta_L \models_{\mathsf{nested}} \alpha[y/x], \alpha[z/x], \Delta'_L, \theta^{\mathsf{IH}}, \Lambda \in E^{\mathsf{IH}} \quad \text{and} \quad \Theta_R \models_{\mathsf{nested}} y \neq_{\mathfrak{U}} z, \Delta'_R, \neg\theta^{\mathsf{IH}}$$

We then have two subcases according to whether $z \in C$ or not

* If $z \in C$, we can take $\theta := \theta^{\mathsf{IH}} \vee y \neq_{\mathfrak{U}} z$ and $E := E^{\mathsf{IH}}$. Their free variables are in C and we only need to check

$$\Theta_L \models_{\mathsf{nested}} \alpha[y/x], \Delta'_L, \theta^{\mathsf{IH}} \vee y \neq_{\mathfrak{U}} z, \Lambda \in E^{\mathsf{IH}} \quad \text{and} \quad \Theta_R \models_{\mathsf{nested}} y \neq_{\mathfrak{U}} z, \Delta'_R, \neg\theta^{\mathsf{IH}} \wedge y =_{\mathfrak{U}} z$$

which follow easily from the induction hypothesis.

* Otherwise, we take $\theta := \theta^{\mathsf{IH}}[y/z]$ and $E := E^{\mathsf{IH}}[y/z]$. In that case, note that we have $\alpha[z/x][y/z] = \alpha[y/x]$ (which would not be necessarily the case if z belonged to C). This allows to conclude that we have

$$\Theta_L \models_{\mathsf{nested}} \alpha[y/x], \Delta'_L, \theta^{\mathsf{IH}}[y/z], \Lambda \in E^{\mathsf{IH}}[y/z] \quad \text{and} \quad \Theta_R \models_{\mathsf{nested}} y \neq_{\mathfrak{U}} z, \Delta'_R, \neg\theta^{\mathsf{IH}}[y/z]$$

directly from the induction hypothesis.

• If the last rule applied is $\times_\eta$, the proof has shape

$$\frac{\Theta_L[\langle x_1, x_2\rangle/x], \Theta_R[\langle x_1, x_2\rangle/x] \vdash \Delta_L[\langle x_1, x_2\rangle/x], \Delta_R[\langle x_1, x_2\rangle/x], \tilde{\mathcal{G}}[\langle x_1, x_2\rangle/x]}{\Theta_L, \Theta_R \vdash \Delta_L, \Delta_R, \tilde{\mathcal{G}}}$$

and one applies the inductive hypothesis as expected to get θ^{IH} and E^{IH} such that

$$\Theta_L[\langle x_1, x_2\rangle/x] \models_{\mathsf{nested}} \Delta_L[\langle x_1, x_2\rangle/x], \theta^{\mathsf{IH}}, \Lambda[\langle x_1, x_2\rangle/x] \in E$$
$$\text{and} \quad \Theta_R[\langle x_1, x_2\rangle/x] \models_{\mathsf{nested}} \Delta_R[\langle x_1, x_2\rangle/x], \neg\theta^{\mathsf{IH}}$$

and with free variables included in C if $x \notin C$ or $(C \cup \{x_1, x_2\}) \setminus \{x\}$ otherwise. In both cases, it is straightforward to check that taking $\theta := \theta^{\mathsf{IH}}[\pi_1(x)/x_1, \pi_2(x)/x_2]$ and $E := E^{\mathsf{IH}}[\pi_1(x)/x_1, \pi_2(x)/x_2]$ will yield the desired result.

• Finally, if the last rule applied is the $\times_\beta$ rule, it has shape

$$\frac{(\Theta_L, \Theta_R)[x_i/x] \vdash (\Delta_L, \Delta_R, \tilde{\mathcal{G}}')[x_i/x]}{(\Theta_L, \Theta_R)[\pi_i(\langle x_1, x_2\rangle)/x] \vdash (\Delta_L, \Delta_R, \tilde{\mathcal{G}}')[\pi_i(\langle x_1, x_2\rangle)/x]}$$

and we can apply the induction hypothesis to get satisfactory θ^{IH} and E^{IH} (moving from $\tilde{\mathcal{G}}$ to $\tilde{\mathcal{G}}'[x_i/x]$ is unproblematic, as we can assume the lemma works for $\tilde{\mathcal{G}}$ with arbitrary subformulas λ and ρ); it is easy to see that we can set $\theta := \theta^{\mathsf{IH}}$ and $E := E^{\mathsf{IH}}$ and conclude.

This completes the proof of Lemma 5.10.

5.3. **Proof of the main result.** We now turn to the proof of our second main result, which we recall from the earlier subsection:

Theorem 5.2 (Effective implicit to explicit for nested data). *Given a witness to the implicit definition of o in terms of $\vec{i}$ relative to Δ_0 $\varphi(\vec{i}, \vec{a}, o)$, one can compute* NRC *expression E such that for any $\vec{i}$, $\vec{a}$ and o, if $\varphi(\vec{i}, \vec{a}, o)$ then $E(\vec{i}) = o$. Furthermore, if the witness is* EL-*normalized, this can be done in polynomial time.*

We have as input a proof of

$$\varphi(\vec{i}, \vec{a}, o) \wedge \varphi(\vec{i}, \vec{a}', o') \to o \equiv_T o'$$

and we want an NRC expression $E(\vec{i})$ such that

$$\varphi(\vec{i}, \vec{a}, o) \models_{\mathsf{nested}} E(\vec{i}) \equiv_T o$$

This will be a consequence of the following theorem.

Theorem 5.11. *Given Δ_0 $\varphi(\vec{i}, \vec{a}, o)$ and $\psi(\vec{i}, \vec{b}, o')$ together with a* EL-*normalized proof with conclusion*

$$\Theta(\vec{i}, \vec{a}, r); \; \varphi(\vec{i}, \vec{a}, r), \psi(\vec{i}, \vec{b}, o') \vdash \exists r' \in_p o'. \, r \equiv_T r'$$

we can compute in polynomial time an NRC *expression $E(\vec{i})$ such that*

$$\Theta(\vec{i}, \vec{a}, r); \; \varphi(\vec{i}, \vec{a}, r), \psi(\vec{i}, \vec{b}, o') \; \models_{\mathsf{nested}} \; r \in E(\vec{i})$$

That is, we can find an NRC query that "collects answers". Assuming Theorem 5.11, let's prove the main result.

Proof of Theorem 5.2. We assume o has a set type, deferring the simple product and Ur-element cases (the latter using GET) to Appendix J. Fix an implicit definition of o up to extensionality relative to $\varphi(\vec{i}, \vec{a}, o)$ and a EL-normalized proof of

$$\varphi(\vec{i}, \vec{a}, o) \; \wedge \; \varphi(\vec{i}, \vec{a}', o') \; \vdash \; o \equiv_{\mathsf{Set}(T)} o'$$

We can apply Lemma 5.7, in the simple case where p is the "empty path", to obtain a EL-normalized derivation of

$$r \in o; \; \varphi(\vec{i}, \vec{a}, o), \varphi(\vec{i}, \vec{a}', o') \vdash \; \exists r' \in o' \; r \equiv_T r' \tag{5.7}$$

Then applying Theorem 5.11 gives an NRC expression $E(\vec{i})$ such that

$$\varphi(\vec{i}, \vec{a}, o) \wedge r \, \hat{\in} \, o \wedge \; \varphi(\vec{i}, \vec{a}', o') \; \models_{\mathsf{nested}} \; r \in E(\vec{i})$$

Thus, the object determined by $\vec{i}$ is always contained in $E(\vec{i})$. Recall that by (5.7), we have a derivation of

$$r \in o; \; \varphi(\vec{i}, \vec{a}, o) \vdash \varphi(\vec{i}, \vec{a}', o') \to \exists r' \in o' \; r \equiv_T r'$$

and applying interpolation (Theorem 5.4) to that gives a Δ_0 formula $\kappa(\vec{i}, r)$ such that the following are valid

$$r \in o \wedge \varphi(\vec{i}, \vec{a}, o) \to \kappa(\vec{i}, r) \tag{5.8}$$

$$\kappa(\vec{i}, r) \wedge \varphi(\vec{i}, \vec{a}', o') \to \exists r' \in o'. \, r \equiv_T r' \tag{5.9}$$

We claim that $E_\kappa(\vec{i}) = \left\{ x \in E(\vec{i}) \mid \kappa(\vec{i}, x) \right\}$ is the desired NRC expression. To show this, assume $\varphi(\vec{i}, \vec{a}, o)$ holds. We know already that $o \subseteq E(\vec{i})$ and, by (5.8), every $r \in o$ satisfies

$\kappa(\vec{i}, o)$, so $o \subseteq E_\kappa(\vec{i})$. Conversely, if $x \in E_\kappa(\vec{i})$, we have $\kappa(\vec{i}, x)$, so by (5.9), we have that $x \in o$, so $E_\kappa(\vec{i}) \subseteq o$. So $E_\kappa(\vec{i}) = o$, which concludes the proof. $\qquad\square$

We now turn to the proof of Theorem 5.11.

Proof of Theorem 5.11. We prove the theorem by induction over the type T. We only prove the inductive step for set types: the inductive case for products is straightforward.

For $T = \mathfrak{U}$, the base case of the induction, it is clear that we can take for E an expression computing the set of all $\mathfrak{U}$-elements in the transitive closure of $\vec{i}$. This can clearly be done in NRC.

So now, we assume $T = \mathsf{Set}(T')$ and that Theorem 5.11 holds up to T'. We have a EL-normalized derivation of

$$\Theta;\ \varphi(\vec{i}, r),\ \psi(\vec{i}, o') \ \vdash\ \exists r' \in_p o'.\ r \equiv_{\mathsf{Set}(T')} r' \tag{5.10}$$

omitting the additional variables for brevity.

From our input derivation, we can easily see that each element of r must be equivalent to some element below o'. This is reflected in Lemma 5.7, which allows us to efficiently compute a proof of

$$\Theta, z \in r;\ \varphi(\vec{i}, r),\ \psi(\vec{i}, o') \ \vdash\ \exists z' \in_{mp} o'.\ z \equiv_{T'} z' \tag{5.11}$$

We can then apply the inductive hypothesis of our main theorem at sort T', which is strictly smaller than $\mathsf{Set}(T')$, on that new proof. This yields an NRC expression $E^{\mathsf{IH}}(\vec{i})$ of type $\mathsf{Set}(T')$ such that

$$\Theta, z \in r;\ \varphi(\vec{i}, r), \psi(\vec{i}, o') \models_{\mathsf{nested}} z \in E^{\mathsf{IH}}(\vec{i})$$

That is, our original hypotheses entail $r \subseteq E^{\mathsf{IH}}(\vec{i})$.

Thus, we have used the inductive hypothesis to get a "superset expression". But now we want an expression that has r as an element. We will do this by unioning a collection of definable subsets of $E^{\mathsf{IH}}(\vec{i})$. To get these, we come back to our input derivation (5.10). By Lemma 5.8, we can efficiently compute a derivation of

$$\Theta;\ \varphi(\vec{i}, r), \psi(\vec{i}, o') \vdash \exists r' \in_p o' \, \forall z \in a\ (z \,\hat{\in}\, r \leftrightarrow z \,\hat{\in}\, r')$$

where we take a to be a fresh variable of sort $\mathsf{Set}(T')$. Now, applying our NRC Parameter Collection result (Theorem 5.9) we obtain an NRC expression $E^{\mathrm{coll}}(\vec{i}, a)$ satisfying

$$\Theta;\ \varphi(\vec{i}, r), \psi(\vec{i}, o') \ \models_{\mathsf{nested}} a \cap r \in E^{\mathrm{coll}}(\vec{i}, a)$$

Now, recalling that we have $r \subseteq E^{\mathsf{IH}}(\vec{i})$ and instantiating a to be $E^{\mathsf{IH}}(\vec{i})$, we can conclude that

$$\Theta;\ \varphi(\vec{i}, r), \psi(\vec{i}, o') \ \models_{\mathsf{nested}} r \in E^{\mathrm{coll}}(\vec{i}, E^{\mathsf{IH}}(\vec{i}))$$

Thus we can take $E^{\mathrm{coll}}(\vec{i}, E^{\mathsf{IH}}(\vec{i}))$ as an explicit definition.

Complexity. Now let us sketch the complexity analysis of the underlying transformation. The main induction is the type T of the object to be defined, and most of the lemmas we use have a complexity that depend on the size of the proofs, which is commensurate with the size of the proof tree multiplied by the size of the input sequent, that we will write n. Let us call $C(T, n)$ the time complexity of our procedure and show it can be taken to be polynomial. For the base case and the product case, we have that $C(\mathfrak{U}, n)$ is $\mathcal{O}(n^k)$ for $k \geq 1$. For set types $\mathsf{Set}(T)$, we first have a polynomial-time procedure in n to obtain the new proof in Lemma 5.11, and we have a recursive call on this proof. Note that Lemma 5.7 also tells us that this proof has size at most n, so the recursive call has complexity at most $C(T, n)$. Then the subsequent transformations are simply polynomial-time on the input proof, so we have for some exponent k large enough

$$C(\mathsf{Set}(T), n) \leq C(T, n) + n^k$$

A similar analysis for products yields that

$$C(T_1 \times T_2, n) \leq C(T_1, n) + C(T_2, n) + n^k$$

All in all, if we call $s(T)$ the size of a type defined in the obvious way, we have

$$C(T, n) = \mathcal{O}(s(T)n^k)$$

by induction on T, so since $s(T) \leq n$, the overall time complexity is indeed polynomial in the size of the input derivation. $\square$

6. Discussion and future work

Our first contribution implies that whenever a set of NRC views determines an NRC query, the query is rewritable over the views in NRC. By our second result, from a proof witnessing determinacy in our EL-normalized proof system, we can efficiently generate the rewriting. Both results apply to a setting where we have determinacy with respect to constraints and views, as in Example 1.2, or to general Δ_0 implicit definitions that may not stem from views.

In terms of impact on databases, a crucial limitation of our work is that we do not yet know how to find the proofs. In the case of relational data, we know of many "islands of decidability" where proofs of determinacy can be found effectively – e.g. for views and queries in guarded logics [BBtC18]. But it remains open to find similar decidability results for views/queries in fragments of NRC.

It is possible to use our proof system without full automation – simply search for a proof, and then when one finds one, generate the rewriting. We have had some success with this approach in the relational setting, where standard theorem proving technology can be applied [BKMT17]. But for the proof systems proposed here, we do not have either our own theorem prover or a reduction to a system that has been implemented in the past. The need to find proofs automatically is pressing since our system is so low-level that it is difficult to do proofs by hand. Indeed, a formal proof of implicit definability for Example 1.2, or even the simpler Example 5.1, would come to several pages.

In [BP21], we introduce an intuitionistic version of our proof system, and give a specialized algorithm for generating NRC transformations from proofs of implicit definability within this system. The algorithm for the intuitionistic case is considerably simpler than for the proof systems we present here, and would probably make a good starting point for an implementation of the system.

The implicit-to-explicit methodology requires a proof of implicit definability, which implies implicit definability over all instances, not just finite ones. This requirement is necessary: one cannot hope to convert implicit definitions over finite instances to explicit NRC queries, even ineffectively. We do not believe that this is a limitation in practice. See Appendix A for details.

For the effective Beth result, the key proof tool was the NRC Parameter Collection theorem, Theorem 5.9. There is an intuition behind this theorem that concerns a general setting, where we have a first-order theory Σ that factors into a conjunction of two formulas $\Sigma_L \wedge \Sigma_R$, and from this we have a notion of a "left formula" (with predicates from Σ_L), a "right formula" (predicates from Σ_R), and a "common formula" (all predicates occur in both Σ_L and Σ_R). Under the hypothesis that a left formula λ is definable from a right formula with parameters, we can conclude that the left formula must actually be definable from a common formula with parameters: see Appendix K for a formal version and the corresponding proof.

Our work contributes to the broader topic of proof-theoretic vs model-theoretic techniques for interpolation and definability theorems. For Beth's theorem, there are reasonably short model-theoretic [Lyn59, CK92] and proof-theoretic arguments [Cra57b, Fit96]. In database terms, you can argue semantically that relational algebra is complete for rewritings of queries determined by views, and producing a rewriting from a proof of determinacy is not that difficult. But for a number of results on definability proved in the 60's and 70's [Cha64, Mak64, Kue71, Gai74], there are short model-theoretic arguments, but no proof-theoretic ones. For our NRC analog of Beth's theorem, the situation is more similar to the latter case: the model-theoretic proof of completeness is relatively short and elementary, but generating explicit definitions from proofs is much more challenging. We hope that our results and tools represent a step towards providing effective versions, and towards understanding the relationship between model-theoretic and proof-theoretic arguments.

References

[AMN08]   H. Andréka, J. X. Madarász, and I. Németi. Definability of new universes in many-sorted logic, 2008. manuscript available at `old.renyi.hu/pub/algebraic-logic/kurzus10/amn-defi.pdf`.

[BBtC18]   Vince Bárány, Michael Benedikt, and Balder ten Cate. Some model theory of guarded negation. *J. Symb. Log.*, 83(4):1307–1344, 2018.

[BBV19]   Michael Benedikt, Pierre Bourhis, and Michael Vanden Boom. Definability and interpolation within decidable fixpoint logics. *Log. Methods Comput. Sci.*, 15(3):29:1–29:53, 2019.

[BCLT16]   Michael Benedikt, Balden Ten Cate, Julien Leblay, and Efthymia Tsamoura. *Generating Plans from Proofs: The Interpolation-Based Approach to Query Reformulation.* Morgan Claypool, San Rafael, CA, 2016.

[BDK18]   Mikolaj Bojanczyk, Laure Daviaud, and Shankara Narayanan Krishna. Regular and first-order list functions. In *LICS*, 2018.

[Bet53]   E. W. Beth. On Padoa's method in the theory of definitions. *Indag. Mathematicae*, 15:330 – 339, 1953.

[BK09]   Michael Benedikt and Christoph Koch. From XQuery to Relational Logics. *ACM TODS*, 34(4):25:1–25:48, 2009.

[BKMT17]   Michael Benedikt, Egor V. Kostylev, Fabio Mogavero, and Efthymia Tsamoura. Reformulating queries: Theory and practice. In *IJCAI*, 2017.

[BNTW95]   Peter Buneman, Shamim A. Naqvi, Val Tannen, and Limsoon Wong. Principles of programming with complex objects and collection types. *Theor. Comput. Sci.*, 149(1):3–48, 1995.

[BP21]   Michael Benedikt and Cécilia Pradic. Generating collection transformations from proofs. In *POPL*, 2021.

[BPW23]   Michael Benedikt, Cécilia Pradic, and Christoph Wernhard. Synthesizing nested relational queries from implicit specifications. In *PODS*, pages 33–45, 2023.

[BtCV16]   Michael Benedikt, Balder ten Cate, and Michael Vanden Boom. Effective interpolation and preservation in guarded logics. *ACM TOCL*, 17(2):8:1–8:46, 2016.

[CH99]   Zoé Chatzadakis and Ehud Hrushovski. Model theory of difference fields. *Transactions of the American Mathematical Society*, 351:2997–3071, 1999.

[Cha64]   C. C. Chang. Some new results in definability. *Bull. of the AMS*, 70(6):808 – 813, 1964.

[CK92]   C. C. Chang and H. Jerome Keisler. *Model Theory*. North-Holland, 1992.

[CL07]   Thomas Colcombet and Christof Löding. Transforming structures by set interpretations. *Logical Methods in Computer Science*, 3(2), 2007.

[Cra57a]   William Craig. Linear reasoning. a new form of the Herbrand-Gentzen theorem. *J. Symb. Log.*, 22(03):250–268, 1957.

[Cra57b]   William Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *J. Symb. Log.*, 22(3):269–285, 1957.

[DH00]   Giovanna D'Agostino and Marco Hollenberg. Logical Questions Concerning The mu-Calculus: Interpolation, Lyndon and Los-Tarski. *J. Symb. Log.*, 65(1):310–332, 2000.

[Fit96]   Melvin Fitting. *First-order Logic and Automated Theorem Proving*. Springer, second edition, 1996.

[FKN13]   Enrico Franconi, Volha Kerhet, and Nhung Ngo. Exact query reformulation over databases with first-order and description logics ontologies. *J. Artif. Int. Res.*, 48:885–922, 2013.

[Gai74]   Haim Gaifman. Operations on relational structures, functors and classes I. In *Proc. of the Tarski Symposium*, volume 25 of *Proc. of Symposia in Pure Mathematics*, pages 20–40, 1974.

[GHHW18]   Jeremy Gibbons, Fritz Henglein, Ralf Hinze, and Nicolas Wu. Relational algebra by way of adjunctions. *PACMPL*, 2(ICFP), 2018.

[Gib16]   Jeremy Gibbons. Comprehending Ringads - for Phil Wadler, on the occasion of his 60th birthday. In *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*, 2016.

[HHM90]   Wilfrid Hodges, I.M. Hodkinson, and Dugald Macpherson. Omega-categoricity, relative categoricity and coordinatisation. *Annals of Pure and Applied Logic*, 46(2):169 – 199, 1990.

[HMO99]   Eva Hoogland, Maarten Marx, and Martin Otto. Beth definability for the guarded fragment. In *LPAR*, 1999.

[Hod75]   Wilfrid Hodges. A normal form for algebraic constructions II. *Logique et Analyse*, 18(71/72):429–487, 1975.

[Hod93]   Wilfrid Hodges. *Model Theory*. Cambridge University Press, 1993.

[Hru14]   Ehud Hrushovski. Groupoids, imaginaries and internal covers. *Turkish Journal of Mathematics*, 36:173 – 198, 2014.

[Hua95]   Guoxiang Huang. Constructing Craig interpolation formulas. In *Computing and Combinatorics*. 1995.

[Jec03]   Thomas Jech. *Set Theory*. Springer, 2003.

[Kle52]   S. C. Kleene. Permutability of inferences in Gentzen's calculi lk and lj. *Memoirs of the American Mathematical Society*, 10:1–26, 1952.

[Koc06]   Christoph Koch. On the Complexity of Non-recursive XQuery and Functional Query Languages on Complex Values. *ACM TODS*, 31(4):1215–1256, 2006.

[Kol90]   Phokion G. Kolaitis. Implicit definability on finite structures and unambiguous computations. In *LICS*, 1990.

[Kue71]   David Kueker. Generalized interpolation and definability. *Annals of Mathematical Logic*, 1(4):423–468, 1971.

[LE65]   E. G. K. Lopez-Escobar. An interpolation theorem for denumerably long sentences. *Fundamenta Mathametica*, 57:253–272, 1965.

[Lyn59]   Roger C. Lyndon. An interpolation theorem in the predicate calculus. *Pacific J. Math.*, 9:129–142, 1959.

[Mak64]   Michael Makkai. On a generalization of a theorem of E. W. Beth. *Acta Math. Ac. Sci. Hung.*, 15:227–235, 1964.

[MBB06]   Erik Meijer, Brian Beckman, and Gavin Bierman. LINQ: Reconciling object, relations and XML in the .NET framework. In *SIGMOD*, 2006.

[MGL+10]   Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer, Shiva Shivakumar, Matt Tolton, and Theo Vassilakis. Dremel: Interactive Analysis of Web-Scale Datasets. *PVLDB*, 3(1-2):330–339, 2010.

[Mos49]   Andrzej Mostowski. An undecidable arithmetical statement. *Fundamenta Mathematicae*, 36(1):143–164, 1949.

[NSV10]   Alan Nash, Luc Segoufin, and Victor Vianu. Views and queries: Determinacy and rewriting. *ACM TODS*, 35(3):1–41, 2010.

[NvP98]   Sara Negri and Jan von Plato. Cut elimination in the presence of axioms. *Bull. Symb. Log.*, 4(4):418–435, 1998.

[NvP01]   Sara Negri and Jan von Plato. *Structural Proof Theory*. Cambridge University Press, 2001.

[Ott00]   Martin Otto. An interpolation theorem. *Bull. Symb. Log.*, 6(4):447–462, 2000.

[Smu68a]   Raymond Smullyan. *First Order Logic*. Springer, 1968.

[Smu68b]   Raymond M. Smullyan. *Craig's Interpolation Lemma and Beth's Definability Theorem.* In: *First-Order Logic*, pages 127–133. Springer, 1968.

[Suc95]   Dan Suciu. *Parallel Programming Languages for Collections*. PhD thesis, Univ. Pennsylvania, 1995.

[SV05]   Luc Segoufin and Victor Vianu. Views and queries: Determinacy and rewriting. In *PODS*, 2005.

[Tak87]   Gaisi Takeuti. *Proof Theory*. North-Holland, second edition, 1987.

[tCFS13]   Balder ten Cate, Enrico Franconi, and Inanç Seylan. Beth definability in expressive description logics. *J. Artif. Int. Res.*, 48(1):347–414, 2013.

[TS00]   Arne S. Troelstra and Helmut Schwichtenberg. *Basic Proof Theory*. Cambridge University Press, 2000.

[TW11]   David Toman and Grant Weddell. *Fundamentals of Physical Design and Query Compilation*. Morgan Claypool, 2011.

[Van01]   Jan Van den Bussche. Simulation of the Nested Relational Algebra by the Flat Relational Algebra, with an Application to the Complexity of Evaluating Powerset Algebra Expressions. *Theor. Comput. Sci.*, 254(1–2):363–377, 2001.

[Won94]   Limsoon Wong. *Querying Nested Collections*. PhD thesis, Univ. Pennsylvania, 1994.

Appendix A. Comparison to the situation with finite instances

Our result concerns a specification $\Sigma(\vec{I}, O \ldots)$ such that $\vec{I}$ implicitly defines O. This can be defined "syntactically" – via the existence of a proof (e.g. in our own proof system). Thus, the class of queries that we deal with could be called the "provably implicitly definable queries". The same class of queries can also be defined semantically, and this is how implicitly defined queries are often presented. But in order to be equivalent to the proof-theoretic version, we need the implicit definability of the object O over $\vec{I}$ to holds considering all nested relations $\vec{I}, O \ldots$, not just finite ones. Of course, the fact that when you phrase the property semantically requires referencing unrestricted instance does not mean that our results depend on the existence of infinite nested relations.

Discussions of finite vs. unrestricted instances appear in many other papers (e.g. [BP21]). And the results in this submission do not raise any new issues with regard to the topic. But we discuss what happens if we take the obvious analog of the semantic definition, but using only finite instances. Let us say that a Δ_0 specification $\Sigma(\vec{I}, O\vec{A})$ *implicitly defines* O *in terms of* $\vec{I}$ *over finite instances* if for any *finite* nested relations $\vec{I}, O, \vec{A}, O', \vec{A}'$, if $\Sigma(\vec{I}, O, \vec{A}) \wedge \Sigma(\vec{I}, O', \vec{A}')$ holds, then $O = O'$. If this holds, then Σ defines a query, and we call such a query *finitely implicitly definable*.

This class of queries is reasonably well understood, and we summarize what is known about it:

- *Can finitely implicitly definable queries always be defined in* NRC*?* The answer is a resounding "no": one can implicitly define the powerset query over finite nested relations. Bootstrapping this, one can define iterated powersets, and show that the expressiveness of implicit definitions is the same as queries in NRC enhanced with powerset – a query language with non-elementary complexity. Even in the setting of relational queries, considering only finite instances leads to a query class that is not known to be in PTIME [Kol90].
- *Can we generate explicit definitions from specifications* Σ, *given a proof that* Σ *implicitly defines* O *in terms of* $\vec{I}$ *over finite instances?* It depends on what you mean by "a proof", but in some sense there is no way to make sense of the question: there is no recursively enumerable complete proof system for such definitions. This follows from the fact that the set of finitely implicitly definable queries is not computably enumerable.
- *Is sticking to specifications* Σ *that are implicit definitions over all inputs – as we do in this work – too strong?* Here the answer can not be definitive. But we know of no evidence that this is too restrictive in practice. Implicit specifications suffice to specify any NRC query. And the answer to the first question above says that if we modified the definition in the obvious way to get a larger class, we would allow specification of queries that do not admit efficient evaluation. The answer to the second question above says that we do not have a witness to membership in this larger class.

Appendix B. Proof of Proposition 3.5: obtaining NRC expressions that verify Δ_0 formulas

Recall that in the body of the paper, we claimed the following statement, concerning the equivalence of NRC expressions of Boolean type and Δ_0 formulas:

There is a polynomial time function taking an extended Δ_0 formula $\varphi(\vec{x})$ and producing an NRC expression $\mathsf{Verify}_\varphi(\vec{x})$, where the expression takes as input $\vec{x}$ and returns true if and only if φ holds.

We refer to this as the "Verification Proposition" later on in these supplementary materials.

Proof. First, one should note that every term in the logic can be translated to a suitable NRC expression of the same sort. For example, a variable in the logic corresponds to a variable in NRC.

We prove the proposition by induction over the formula $\varphi(\vec{x})$; we only treat the case for half of the connectives, as the dual connectives can be then be handled similarly using De Morgan rules and the fact that boolean negation is definable in NRC:

- If $\varphi(\vec{x})$ is an equality $t =_T t'$, then one can translate that to the NRC expression $\bigcup\{\{\langle\rangle\} \mid z \in \{t\} \cap \{t'\}\}$.
- If $\varphi(\vec{x})$ is a membership $t \in_T t'$, then it can be translated to the NRC expression $\bigcup\{\{\langle\rangle\} \mid z \in \{\{t\}\} \cap \{t'\}\}$.
- If $\varphi(\vec{x})$ is a disjunction $\varphi_1(\vec{x}) \vee \varphi_2(\vec{x})$, we take $\mathsf{Verify}_\varphi(\vec{x}) = \mathsf{Verify}_{\varphi_1}(\vec{x}) \cup \mathsf{Verify}_{\varphi_2}(\vec{x})$.
- If $\varphi(\vec{x})$ begins with a bounded existential quantification $\exists z \in y\ \psi(\vec{x}, y, z)$, we simply set $\mathsf{Verify}_\varphi(\vec{x}, y) = \bigcup\{\mathsf{Verify}_{\psi(\vec{x}, y, z)} \mid z \in y\}$. $\qquad\square$

Note that the converse (without the polynomial time bound) also holds; this will follow from the more general result on moving from NRC to interpretations that is proven later in the supplementary materials.

Appendix C. First part of Proposition 4.10: Reduction to monadic schemas for NRC

In the body of the paper we mentioned that it is possible to reduce questions about definability within NRC to the case of monadic schemas. We now give the details of this reduction.

Recall that *monadic type* is a type built only using the atomic type $\mathfrak{U}$ and the type constructor Set. Monadic types are in one-to-one correspondence with natural numbers by setting $\mathfrak{U}_0 := \mathfrak{U}$ and $\mathfrak{U}_{n+1} := \mathsf{Set}(\mathfrak{U}_n)$. A monadic type is thus a $\mathfrak{U}_n$ for some $n \in \mathbb{N}$. A nested relational schema is monadic if it contains only monadic types, and a Δ_0 formula is said to be monadic if it all of its variables have monadic types.

We start with a version of the reduction only for NRC expressions:

Proposition C.1. *For any nested relational schema $\mathcal{SCH}$, there is a monadic nested relational schema $\mathcal{SCH}'$, an injection Convert from instances of $\mathcal{SCH}$ to instances of $\mathcal{SCH}'$ that is definable in NRC, and an NRC[GET] expression Convert^{-1} such that Convert^{-1}∘Convert is the identity transformation from $\mathcal{SCH} \to \mathcal{SCH}$.*

Furthermore, there is a Δ_0 formula $\mathsf{Im}_{\mathsf{Convert}}$ from $\mathcal{SCH}'$ to Bool such that $\mathsf{Im}_{\mathsf{Convert}}(i')$ holds if and only if $i' = \mathsf{Convert}(i)$ for some instance i of $\mathcal{SCH}$.

To prove this we give an encoding of general nested relational schemas into monadic nested relational schemas that will allow us to reduce the equivalence between NRC expression, interpretations, and implicit definitions to the case where input and outputs are monadic.

Note that it will turn out to be crucial to check that this encoding may be defined *either* through NRC expressions or interpretations, but in this subsection we will give the definitions in terms of NRC expressions.

The first step toward defining these encodings is actually to emulate in a sound way the cartesian product structure for types $\mathfrak{U}_n$. Here "sound" means that we should give terms for pairing and projections that satisfy the usual equations associated with cartesian product structure.

Proposition C.2. *For every $n_1, n_2 \in \mathbb{N}$, there are NRC expressions $\widehat{\mathsf{Pair}}(x, y) : \mathfrak{U}_{n_1}, \mathfrak{U}_{n_2} \to \mathfrak{U}_{\max(n_1, n_2)+2}$ and NRC expressions $\widehat{\pi}_i(x) : \mathfrak{U}_{\max(n_1, n_2)+2} \to \mathfrak{U}_{n_i}$ for $i \in \{1, 2\}$ such that the following equations hold*

$$\widehat{\pi}_1\left(\widehat{\mathsf{Pair}}(a_1, a_2)\right) = a_1 \qquad\qquad \widehat{\pi}_2\left(\widehat{\mathsf{Pair}}(a_1, a_2)\right) = a_2$$

Furthermore, there is a Δ_0 formula $\mathsf{Im}_{\widehat{\mathsf{Pair}}}(x)$ such that $\mathsf{Im}_{\widehat{\mathsf{Pair}}}(a)$ holds if and only if there exists a_1, a_2 such that $\widehat{\mathsf{Pair}}(a_1, a_2) = a$. In such a case, the following also holds

$$\widehat{\mathsf{Pair}}(\widehat{\pi}_1(a), \widehat{\pi}_2(a)) \;=\; a$$

Proof. We adapt the Kuratowski encoding of pairs $(a, b) \mapsto \{\{a\}, \{a, b\}\}$. The notable thing here is that, for this encoding to make sense in the typed monadic setting, the types of a and b need to be the same. This will not be an issue because we have NRC-definable embeddings

$$\uparrow_n^m \colon \mathfrak{U}_n \to \mathfrak{U}_m$$

for $n \leq m$ defined as the $m - n$-fold composition of the singleton transformation $x \mapsto \{x\}$. This will be sufficient to define the analogues of pairing for monadic types and thus to define $\mathsf{Convert}_T$ by induction over T. On the other hand, $\mathsf{Convert}_T^{-1}$ will require a suitable encoding of projections. This means that to decode an encoding of a pair, we need to make use of a transformation inverse to the singleton construct $\uparrow$. But we have this thanks to the GET construct. We let

$$\downarrow_n^m \colon \mathfrak{U}_m \to \mathfrak{U}_n$$

the transformation inverse to $\uparrow_n^m$, defined as the $m - n$-fold composition of GET.

    Firstly, we define the family of transformations $\widehat{\mathsf{Pair}}_{n_1, n_2}(x_1, x_2)$, where x_i is an input of type $\mathfrak{U}_{n_i}$ for $i \in \{1, 2\}$ and the output is of type $\mathfrak{U}_{\max(n_1, n_2)+2}$, as follows

$$\widehat{\mathsf{Pair}}_{n_1, n_2}(x_1, x_2) \;:=\; \{\{\uparrow x_1\}, \{\uparrow x_1, \uparrow x_2\}\}$$

    The associated projections $\widehat{\pi}_i^{n_1, n_2}(x)$ where x has type $\mathfrak{U}_{\max(n_1, n_2)+2}$ and the output is of type $\mathfrak{U}_{n_i}$ are a bit more challenging to construct. The basic idea is that there is first a case distinction to be made for encodings $\widehat{\mathsf{Pair}}_{n_1, n_2}(x_1, x_2)$: depending on whether $\uparrow x_1 = \uparrow x_2$ or not. This can be actually tested by an NRC expression. Once this case distinction is made, one may informally compute the projections as follows:

- if $\uparrow x_1 = \uparrow x_2$, both projections can be computed as a suitable downcasting $\downarrow$ (the depth of the downcasting is determined by the output type, which is not necessarily the same for both projections).
- otherwise, one needs to single out the singleton $\{\uparrow x_1\}$ and the two-element set $\{\uparrow x_1, \uparrow x_2\}$ in NRC. Then, one may compute the first projection by downcasting the singleton, and the second projection by first computing $\{\uparrow x_2\}$ as a set difference and then downcasting with $\downarrow$.

    We now give the formal encoding for projections, making a similar case distinction. To this end, we first define a generic NRC expression

$$\mathsf{AllPairs}_T(x) : \mathsf{Set}(T) \to \mathsf{Set}(T \times T)$$

computing all the pairs of distinct elements of its input x

$$\mathsf{AllPairs}_T(x) = \bigcup \{ \bigcup \{ \{(y, z)\} \mid y \in x \setminus \{z\} \} \mid z \in x \}$$

Note in particular that $\mathsf{AllPairs}(i) = \emptyset$ if and only if i is a singleton or the empty set. The projections can thus be defined as

$$
\begin{aligned}
\widehat{\pi}_1(x) &:= \; \mathsf{case}\,(\mathsf{AllPairs}(x) = \emptyset, \; \downarrow x, \; \downarrow \bigcup \{\pi_1(z) \cap \pi_2(z) \mid z \in \mathsf{AllPairs}(x)\}) \\
\widehat{\pi}_2(x) &:= \; \mathsf{case}\,(\mathsf{AllPairs}(x) = \emptyset, \; \downarrow x, \; \downarrow (x \setminus \uparrow \widehat{\pi}_1(x))))
\end{aligned}
$$

These definitions crucially ensure that, for every object a_i with $i \in \{1, 2\}$, we have

$$\widehat{\pi}_i \left(\widehat{\mathsf{Pair}}(a_1, a_2) \right) = a_i$$

Now all remains to be done is to define $\mathsf{Im}_{\widehat{\mathsf{Pair}}}$. Before that, it is helpful to define a formula $\mathsf{Im}_{\uparrow_n^m}(x)$ which holds if and only if x is in the image of $\mathsf{Im}_{\uparrow_n^m}$.

As a preliminary step, define generic Δ_0 formulas $\mathsf{IsSing}(x)$ and $\mathsf{IsTwo}(x)$ taking an object of type $\mathsf{Set}(T)$ and returning a Boolean indicating whether the object is a singleton or a two-element set. Defining $\mathsf{Im}_{\uparrow_n^m}$ is straightforward using IsSing and Boolean connectives. Then $\mathsf{Im}_{\widehat{\mathsf{Pair}}_{n,n}}(x)$ can be defined as follows for each $n \in \mathbb{N}$

$$
\begin{aligned}
\mathsf{Im}_{\widehat{\mathsf{Pair}}_{n,n}}(x) &:= \left(\mathsf{IsSing}(x) \wedge \mathsf{Im}_{\widehat{\mathsf{Pair}}_{n,n}}^{\mathsf{IsSing}}(x) \right) \vee \left(\mathsf{IsTwo}(x) \wedge \mathsf{Im}_{\widehat{\mathsf{Pair}}_{n,n}}^{\mathsf{IsSing}}(x) \right) \\
\mathsf{Im}_{\widehat{\mathsf{Pair}}_{n,n}}^{\mathsf{IsSing}}(x) &:= \exists z \in x \; \mathsf{IsSing}(z) \\
\mathsf{Im}_{\widehat{\mathsf{Pair}}_{n,n}}^{\mathsf{IsTwo}}(x) &:= \exists z \; z' \in x \; (\mathsf{IsTwo}(z) \wedge \mathsf{IsSing}(z') \wedge \forall y \in z \; y \in z')
\end{aligned}
$$

Then, the more general $\mathsf{Im}_{\widehat{\mathsf{Pair}}_{n_1, n_2}}$ can be defined using $\mathsf{Im}_{\uparrow_{n_i}^m}$ where $m = \max(n_1, n_2)$.

$$\mathsf{Im}_{\widehat{\mathsf{Pair}}_{n_1, n_2}}(x) := \mathsf{Im}_{\widehat{\mathsf{Pair}}_{m,m}}(x) \cap \mathsf{Im}_{\uparrow_{n_1}^m}(\widehat{\pi}_1(x)) \cap \mathsf{Im}_{\uparrow_{n_2}^m}(\widehat{\pi}_2(x))$$

One can then easily check that $\mathsf{Im}_{\widehat{\mathsf{Pair}}}$ does have the advertised property: if $\mathsf{Im}_{\widehat{\mathsf{Pair}}}(a)$ holds for some object a, then there are a_1 and a_2 such that $\widehat{\mathsf{Pair}}(a_1, a_2) = a$ and we have

$$\widehat{\mathsf{Pair}}(\widehat{\pi}_1(a), \widehat{\pi}_2(a)) = a \qquad \qquad \square$$

We are now ready to give the proof of the proposition given at the beginning of this subsection.

Proof. $\mathsf{Convert}_T$, $\mathsf{Convert}_T^{-1}$ and $\mathsf{Im}_{\mathsf{Convert}_T}$ are defined by induction over T. Beforehand, define the map d taking a type T to a natural number $d(T)$ so that $\mathsf{Convert}$ maps instances of type T to monadic types $\mathfrak{U}_{d(T)}$.

$$
\begin{aligned}
d(\mathfrak{U}) &= 0 & d(\mathsf{Set}(T)) &= 1 + d(T) \\
d(T_1 \times T_2) &= 2 + \max(d(T_1), d(T_2)) & d(\mathsf{Unit}) &= 0
\end{aligned}
$$

$\mathsf{Convert}_T$, $\mathsf{Convert}_T^{-1}$ and $\mathsf{Im}_{\mathsf{Convert}_T}$ are then defined by the following rules, where we write $\mathsf{Map}\,(z \mapsto E)\,(x)$ for the NRC expression $\bigcup \{\{E\} \mid z \in x\}$.

$$
\begin{array}{lll}
\mathsf{Convert}_{\mathfrak{U}}(x) & := & x \\
\mathsf{Convert}_{\mathsf{Set}(T)}(x) & := & \mathsf{Map}\,(z \mapsto \mathsf{Convert}_T(z))\,(x) \\
\mathsf{Convert}_{\mathsf{Unit}}(x) & := & c_0 \\
\mathsf{Convert}_{T_1 \times T_2}(x) & := & \widehat{\mathsf{Pair}}(\mathsf{Convert}_{T_1}(\pi_1(x)), \mathsf{Convert}_{T_2}(\pi_2(x)))
\end{array}
$$

$$
\begin{array}{lll}
\mathsf{Convert}_{\mathfrak{U}}^{-1}(x) & := & x \\
\mathsf{Convert}_{\mathsf{Set}(T)}^{-1}(x) & := & \mathsf{Map}\,\big(z \mapsto \mathsf{Convert}_T^{-1}(z)\big)\,(x) \\
\mathsf{Convert}_{\mathsf{Unit}}(x) & := & \langle\rangle \\
\mathsf{Convert}_{T_1 \times T_2}^{-1}(x) & := & \Big\langle \mathsf{Convert}_{T_1}^{-1}(\widehat{\pi}_1(x)), \mathsf{Convert}_{T_2}^{-1}(\widehat{\pi}_2(x)) \Big\rangle
\end{array}
$$

$$
\begin{array}{lll}
\mathsf{Im}_{\mathsf{Convert}_{\mathfrak{U}}}(x) & := & \mathsf{True} \\
\mathsf{Im}_{\mathsf{Convert}_{\mathsf{Set}(T)}}(x) & := & \forall z \in x\ \mathsf{Im}_{\mathsf{Convert}_T}(z) \\
\mathsf{Im}_{\mathsf{Convert}_{T_1 \times T_2}}(x) & := & \mathsf{Im}_{\mathsf{Pair}_{d(T_1),d(T_2)}}(x) \wedge \mathsf{Im}_{\mathsf{Convert}_{T_1}}(\widehat{\pi}_1(x)) \wedge \mathsf{Im}_{\mathsf{Convert}_{T_2}}(\widehat{\pi}_2(x))
\end{array}
$$

It is easy to check, by induction over T, that for every object a of type T

$$
\mathsf{Convert}^{-1}(\mathsf{Convert}(a)) = a
$$

and that for every object b of type $\mathfrak{U}_{d(T)}$, if $\mathsf{Im}_{\mathsf{Convert}_T}(b) = \mathsf{True}$, then it lies in the image of $\mathsf{Convert}_T$ and $\mathsf{Convert}(\mathsf{Convert}^{-1}(b)) = b$. $\qquad\square$

APPENDIX D. SECOND PART OF PROPOSITION 4.10:
MONADIC REDUCTION FOR INTERPRETATIONS

We have seen so far that it is possible to reduce questions about definability within NRC to the case of monadic schema. Now we turn to the analogous statement for interpretations, given by the following proposition:

Proposition D.1. *For any object schema $\mathcal{SCH}$, there is a monadic nested relational schema $\mathcal{SCH}'$, a Δ_0 interpretation $\mathcal{I}_{\mathsf{Convert}}$ from instances of $\mathcal{SCH}$ to instances of $\mathcal{SCH}'$, and another interpretation $\mathcal{I}_{\mathsf{Convert}^{-1}}$ from instances of $\mathcal{SCH}$ to instances of $\mathcal{SCH}'$ compatible with $\mathsf{Convert}$ and $\mathsf{Convert}^{-1}$ as defined in Proposition D.1 in the following sense: for every instance I of $\mathcal{SCH}$ and for every instance J of $\mathcal{SCH}'$ in the codomain of $\mathsf{Convert}$, we have*

$$
\mathsf{Convert}^{-1}(J) = \mathsf{Collapse}(\mathcal{I}_{\mathsf{Convert}^{-1}}(J)) \qquad\qquad \mathsf{Convert}(I) = \mathsf{Collapse}(\mathcal{I}_{\mathsf{Convert}}(I))
$$

Before proving Proposition D.1, it is helpful to check that a number of basic NRC connectives may be defined at the level of interpretations. To do so, we first present a technical result for more general interpretations.

Proposition D.2. *For any sort T, there is an interpretation of $\mathcal{SCH}_T$ into $\mathcal{SCH}_T$ taking a models M whose every sort is non-empty and Bool has at least two elements to a model M of $\mathsf{O}(T)$. Furthermore, we have that M' is (up to isomorphism) the largest quotient of M' satisfying $\mathsf{O}(T)$.*

Proof. This interpretation corresponds to a quotient of the input, that is definable at every sort

$$
\begin{array}{llll}
\varphi_{\equiv}^{\mathsf{Set}(T)}(x,y) & = & \forall z\ (z \in x \leftrightarrow z \in y) & \qquad \varphi_{\equiv}^{T_1 \times T_2}(x,y) = \pi_1(x) = \pi_1(y) \wedge \pi_2(x) = \pi_2(y) \\
\varphi_{\equiv}^{\mathsf{Unit}}(x,y) & = & \top & \qquad \varphi_{\equiv}^{\mathfrak{U}}(x,y) = x =_{\mathfrak{U}} y \qquad\qquad\square
\end{array}
$$

Proposition D.3. *The following Δ_0-interpretations are definable:*
- $\mathcal{I}_{\mathsf{Sing}}$ *defining the transformation* $x \mapsto \{x\}$.
- $\mathcal{I}_\cup$ *defining the transformation* $x, y \mapsto x \cup y$.

Furthermore, assuming that $\mathcal{I}$ is a Δ_0-interpretation defining a transformation E and $\mathcal{I}'$ is a Δ_0-interpretation defining a transformation R, the following Δ_0-interpretations are also definable:
- $\mathsf{Map}(\mathcal{I})$ *defining the transformation* $x \mapsto \{E(y) \mid x \in y\}$.
- $\langle \mathcal{I}, \mathcal{I}' \rangle$ *defining the transformation* $x, y \mapsto (E(x), F(y))$.

Proof.
- For the singleton construction $\{e\}$ with e of type T, we take the interpretation $\mathcal{I}_e$ for e, where e itself is interpreted by a constant c and we add an extra level represented by an input constant c'. Then $\varphi_{\mathsf{Domain}}^{\mathsf{Set}(T)}(x)$ is set to $y = c'$ and $\varphi_\in^T(x, y)$ to $x = c \wedge y = c'$.
- The empty set $\{\}$ at type $\mathsf{Set}(T)$ is given by the trivial interpretation where $\varphi_{\mathsf{Domain}}^{\mathsf{Set}(T)}(x)$ is set to $x = c$ for some constant c and $\varphi_{\mathsf{Domain}}^{T'}$ is set to false for T' a component type of T, as well as all the $\varphi_\in^T$.
- For the binary union $\cup : \mathsf{Set}(T), \mathsf{Set}(T) \to \mathsf{Set}(T)$, the interpretation is easy: T is interpreted as itself. The difference between input and output is that $\mathsf{Set}(T) \times \mathsf{Set}(T)$ is not an output sort and that $\mathsf{Set}(T)$ is interpreted as a single element, the constant $\langle \rangle$ of Unit.

$$
\begin{aligned}
\varphi_{\mathsf{Domain}}^{\mathsf{Set}(T)}(x) &:= x = \langle \rangle \\
\varphi_\in^T(z, x) &:= z \in \pi_1(\mathsf{o}_{in}) \vee z \in \pi_2(\mathsf{o}_{in})
\end{aligned}
$$

- We now discuss the Map operator. Assume that we have an interpretation $\mathcal{I}$ defining a transformation $S \to T$ that we want to lift to an interpretation $\mathsf{Map}(\mathcal{I}) : \mathsf{Set}(S) \to \mathsf{Set}(T)$. Let us write $\psi_{\mathsf{Domain}}^{T'}$, $\psi_\in^{T'}$ and $\psi_\equiv^{T'}$ for the formulas making up $\mathcal{I}$ and reserve the φ formulas for $\mathsf{Map}(\mathcal{I})$. At the level of sort, let us write $\tau^{\mathcal{I}}$ and $\tau^{\mathsf{Map}(\mathcal{I})}$ to distinguish the two.

  For every $T' \leq T$ such that T' is not a cartesian product or a component type of Bool, we set $\tau^{\mathsf{Map}(\mathcal{I})}(T') = S, \tau^{\mathcal{I}}$. This means that objects of sort T' are interpreted as in $\mathcal{I}$ with an additional tag of sort S. We interpret the output object $\mathsf{Set}(T)$ as a singleton by setting $\tau^{\mathsf{Map}(\mathcal{I})}(\mathsf{Set}(T)) = \mathsf{Unit}$.

  Assuming that $T \neq \mathfrak{U}, \mathsf{Unit}$, $\mathsf{Map}(\mathcal{I})$ is determined by setting the following

$$
\begin{aligned}
\varphi_{\mathsf{Domain}}^{\mathfrak{U}}(a) &:= \exists s \in \mathsf{o}_{in} \; \psi_{\mathsf{Domain}}(a)[s/\mathsf{o}_{in}] \\
\varphi_\in^{\mathfrak{U}}(a, s, \vec{x}) &:= \psi_\in^{\mathfrak{U}}(a, \vec{x})[s/\mathsf{o}_{in}] \\[8pt]
\varphi_{\mathsf{Domain}}^{T'}(s, \vec{x}) &:= \psi_{\mathsf{Domain}}^{T'}(x)[s'/\mathsf{o}_{in}] \\
\varphi_\in^{T'}(s, \vec{x}, s', \vec{y}) &:= \exists \vec{x'} \; \psi_\in^{T'}(\vec{x'}, \vec{y})[s'/\mathsf{o}_{in}] \wedge \varphi_\equiv^{T'}(s, \vec{x}, s', \vec{x'}) \\[8pt]
\varphi_{\mathsf{Domain}}^{T}(s, \vec{x}) &:= s \in \mathsf{o}_{in} \\
\varphi_\in^{T}(s, \vec{x}) &:= \varphi_{\mathsf{Domain}}^{T}(s, \vec{x})
\end{aligned}
$$

  where $[x/\mathsf{o}_{in}]$ means that we replace occurrences of the constant o_{in} by the variable x and sorts T' and $T' \times T''$ are component types of T. Note that this definition is technically by induction over the type, as we use $\varphi_\equiv^{T'}$ to define $\varphi_\in^{T'}$. In case T is $\mathfrak{U}$ or Unit, the last two formulas $\varphi_{\mathsf{Domain}}^T$ and $\varphi_\in^T$ need to change. If $T = \mathsf{Unit}$, then we set

$$
\varphi_{\mathsf{Domain}}^{\mathsf{Unit}}(c_0) := \varphi_\in^{\mathsf{Unit}}(c_0, c_0) := \exists s \in \mathsf{o}_{in} \; \top
$$

and if $T = \mathfrak{U}$, we set

$$\varphi_{\mathsf{Domain}}^{\mathfrak{U}}(a) := \varphi_{\in}^{\mathfrak{U}}(a) := \exists s \in \mathsf{o}_{in} \; \psi_{\mathsf{Domain}}(a)[s/\mathsf{o}_{in}]$$

- Finally we need to discuss the pairing of two interpretation-definable transformations $\langle \mathcal{I}_1, \mathcal{I}_2 \rangle : S \to T_1 \times T_2$. Similarly as for map we reserve $\varphi_{\mathsf{Domain}}^T$, $\varphi_{\in}^T$ and $\varphi_{\equiv}^T$ formulas for the interpretation $\langle \mathcal{I}_1, \mathcal{I}_2 \rangle$. We write ψ_{Domain}^T, $\psi_{\in}^T$ and $\psi_{\equiv}^T$ for components of $\mathcal{I}$ and $\theta_{\mathsf{Domain}}^T$, $\theta_{\in}^T$ and $\theta_{\equiv}^T$ for components of $\mathcal{I}'$.

  Now, the basic idea is to interpret output sorts of $\langle \mathcal{I}_1, \mathcal{I}_2 \rangle$ as tagged unions of elements that either come from $\mathcal{I}_1$ or $\mathcal{I}_2$. Here, we exploit the assumption that $\mathcal{SCH}_T$ contains the sort $\mathsf{Bool.}$ and that every sort is non-empty to interpret the tag of the union. The union itself is then encoded as a concatenation of a tuple representing a would-be element form $\mathcal{I}_1$ with another tuple representing a would-be element from $\mathcal{I}_2$, the correct component being selected with the tag. For that second trick to work, note that we exploit the fact that every sort has a non-empty denotation in the input structure. Concretely, for every T component type of either T_1 or T_2, we thus set

$$
\begin{aligned}
\tau^{\langle \mathcal{I}_1, \mathcal{I}_2 \rangle}(T) \quad &:= \quad \mathsf{Bool}, \tau^{\mathcal{I}_1}(T), \tau^{\mathcal{I}_2}(T) \\
\varphi_{\mathsf{Domain}}^T(u, \vec{x}, \vec{y}) \quad &:= \quad (u = \mathsf{tt} \wedge \psi_{\mathsf{Domain}}^T(\vec{x})) \vee (u \neq \mathsf{tt} \wedge \theta_{\mathsf{Domain}}^T(\vec{y})) \\
\varphi_{\in}^T(u, \vec{x}, \vec{y}, u', \vec{x}', \vec{y}') \quad &:= \quad (u = u' = \mathsf{tt} \wedge \psi_{\in}^T(\vec{x}, \vec{x}')) \vee (u = u' = \mathsf{ff} \wedge \theta_{\in}^T(\vec{y}, \vec{y}')) \\
\varphi_{\equiv}^T(u, \vec{x}, \vec{y}, u', \vec{x}', \vec{y}') \quad &:= \quad (u = u' = \mathsf{tt} \wedge \psi_{\equiv}^T(\vec{x}, \vec{x}')) \vee (u = u' = \mathsf{ff} \wedge \theta_{\equiv}^T(\vec{y}, \vec{y}'))
\end{aligned}
$$

  Note that this interpretation does not quite correspond to a pairing because it is not a complex object interpretation: the interpretation of common subobjects of T_1 and T_2 are not necessarily identified, so the output is not necessarily a model of O. This is fixed by postcomposing with the interpretation of Proposition D.2 to obtain $\langle I_1, I_2 \rangle$. $\qquad\square$

Proof of Proposition D.1. Similarly as with Proposition C.1, we define auxiliary interpretations $\mathcal{I}_\uparrow$, $\mathcal{I}_\downarrow$ $\mathcal{I}_{\widehat{\mathsf{Pair}}}$, $\mathcal{I}_{\hat{\pi}_1}$ and $\mathcal{I}_{\hat{\pi}_2}$ mimicking the relevant constructs of Proposition C.1. Then we will dispense with giving the recursive definitions of $\mathcal{I}_{\mathsf{Convert}_T}$ and $\mathcal{I}_{\mathsf{Convert}_T^{-1}}$, as they will be obvious from inspecting the cases given in the proof of Proposition C.1 and replicating them using Proposition D.3 together with closure under composition of interpretations.

$\mathcal{I}_\uparrow$, $\mathcal{I}_\downarrow$ and $\mathcal{I}_{\mathsf{Pair}}$ are easy to define through Proposition D.3, so we focus on the projections $\mathcal{I}_{\hat{\pi}_1^{n_1, n_2}}$ and $\mathcal{I}_{\hat{\pi}_2^{n_1, n_2}}$, defining transformations from $\mathfrak{U}_m$ to $\mathfrak{U}_{n_i}$ for $i \in \{1, 2\}$ where $m := \max(n_1, n_2)$. Note that in both cases, the output sort is part of the input sorts. Thus an output sort will be interpreted by itself in the input, and the formulas will be trivial for every sort lying strictly below the output sort: we take

$$\varphi_{\in_{\mathfrak{U}_k}}(x, y) \; := \; x \in y \wedge \varphi_{\mathsf{Domain}}^{\mathfrak{U}_{k+1}}(y) \qquad \varphi_{\equiv}^{\mathfrak{U}_k}(x, y) \; := \; x = y \qquad \varphi_{\mathsf{Domain}}^{\mathfrak{U}_k}(x) \; := \; \top$$

for every $k < n_i$ (i according to which projection we are defining). The only remaining important data that we need to provide are the formulas $\varphi_{\mathsf{Domain}}^{\mathfrak{U}_{n_i}}$, which, of course, differ for both projections. We provide those below, calling o_{in} the designated input object. For both cases, we use an auxiliary predicate $x \in^k y$ standing for $\exists y_1 \in y \ldots \exists y_{k-1} \in y_{k-2} \; x \in y_{k-1}$ for $k > 1$; for $k = 0, 1$, we take $x \in^1 y$ to be $x \in y$ and $x \in^0 y$ for $x = y$.

- For $\mathcal{I}_{\hat{\pi}_1^{n_1, n_2}}$, we set

$$\varphi_{\mathsf{Domain}}^{\mathfrak{U}_{n_1}}(x) \; := \; \forall z \in o_{in} \; \exists z' \in z \; x \in^{m-n_1} z'$$

The basic idea is that the outermost $\forall\exists$ ensures that we compute the intersection of the two sets contained in the encoding of the pair.

- For $\mathcal{I}_{\hat{\pi}_2^{n_1,n_2}}$, first note that there are obvious Δ_0-predicates $\mathsf{IsSing}(x)$ and $\mathsf{IsTwo}(x)$ classifying singletons and two element sets. This allows us to write the following Δ_0 formula

$$\varphi_{\mathsf{Domain}}^{\mathfrak{U}_{n_1}}(x) \; := \; \bigvee \left[\begin{array}{l} \mathsf{IsSing}(x) \wedge \forall z \in o_{in} \; \exists z' \in z \; x \in^{m-n_2} z' \\ \mathsf{IsTwo}(x) \wedge \exists z \; z' \in o_{in} \; \exists y \in z' \; (y \notin z \wedge x \in^{m-n_2} z') \end{array} \right.$$

It is then easy to check that, regarded as transformations, those interpretation also implement the projections for Kuratowski pairs.　　　　　　　　　　　　　　　　　□

Appendix E. Proof of Theorem 4.8: converting between NRC expressions and interpretations

In the body of the paper we claimed that NRC expressions have the same expressiveness as interpretations. One direction of this expressive equivalence is given in the following lemma:

Lemma E.1. *There is an* EXPTIME *computable function taking an* NRC *expression E to an equivalent FO interpretation $\mathcal{I}_E$.*

As we mentioned in the body of the paper, very similar results occur in the prior literature, going as far back as [Van01].

Proof. We can assume that the input and output schemas are monadic, using the reductions to monadic schemas given previously. Indeed, if we solve the problem for expressions where input and output schemas are monadic, we can reduce the problem of finding an interpretation for an arbitrary NRC expression $E(x)$ as follows: construct a Δ_0 interpretation $\mathcal{I}$ for the expression $\mathsf{Convert}(E(\mathsf{Convert}^{-1}(x)))$ – where $\mathsf{Convert}$ and $\mathsf{Convert}^{-1}$ are taken as in Proposition C.1 – and then, using closure under composition of interpretations (see e.g. [BK09]), one can then leverage Proposition D.1 to produce the composition of $\mathcal{I}_{\mathsf{Convert}^{-1}}$, $\mathcal{I}$ and $\mathcal{I}_{\mathsf{Convert}}$ which is equivalent to the original expression E.

The argument proceeds by induction on the structure of $E : \vec{T} \to S$ in NRC. Some atomic operators were treated in the prior section, like singleton $\cup$, tupling, and projections. Using closure of interpretations under composition, we are thus able to translate compositions of those operators. We are only left with a few cases.

- For the set difference, since interpretations are closed under composition, it suffices to prove that we can code the transformation

$$(x, y) \; \mapsto \; x \setminus y$$

at every sort $\mathsf{Set}(\mathfrak{U}_n)$. Each sort gets interpreted by itself. We thus set

$$\begin{array}{lll} \varphi_{\mathsf{Domain}}^{\mathfrak{U}_n}(z) & := & z \in \pi_1(o_{in}) \wedge z \notin \pi_1(o_{in}) \\ \varphi_{\mathsf{Domain}}^{\mathfrak{U}_k}(z) & := & \exists z' \; (\varphi_{\mathsf{Domain}}^{\mathfrak{U}_n} \wedge z \in^{n-k} z') \\ \varphi_{\in}^{\mathfrak{U}_k}(z, z') & := & z \in z' \wedge \varphi_{\mathsf{Domain}}^{\mathfrak{U}_k}(z) \wedge \varphi_{\mathsf{Domain}}^{\mathfrak{U}_{k+1}}(z') \end{array}$$

- To get NRC expressions, it suffices to create a Δ_0 interpretation corresponding to GET which follows

$$\varphi_{\mathsf{Domain}}^{\mathfrak{U}}(a) \; := \; (\exists! \; z \in o_{in} \; z = a) \vee (\neg(\exists! \; z \in o_{in}) \wedge a = c_0)$$

- For the binding operator

$$\bigcup\{E_1 \mid x \in E_2\}$$

we exploit the classical decomposition

$$\bigcup \circ \, \mathsf{Map}(E_1) \, \circ \, E_2$$

As interpretations are closed under composition and the mapping operations was handled in Proposition D.3, it suffices to give an interpretation for the expression $\bigcup : \mathsf{Set}(\mathsf{Set}(T)) \to \mathsf{Set}(T)$ for every sort T. This is straightforward: each sort gets interpreted as itself, except for $\mathsf{Set}(T)$ itself which gets interpreted as the singleton $\{c_0\}$. The only non-trivial case is the following

$$\varphi_\in^T(x,y) \; := \; \varphi_\mathsf{Domain}^T \; := \; \exists y' \in \mathsf{o}_{in} \; x \in y' \qquad\qquad \square$$

From interpretations to NRC expressions. The other direction of the expressive equivalence is provided by the following lemma:

Lemma E.2. *There is a polynomial time function taking a Δ_0 interpretation to an equivalent NRC expression.*

This direction is not used directly in the conversion from implicitly definable transformations to NRC, but it is of interest in showing that NRC and Δ_0 interpretations are equally expressive.

Proof of Lemma E.2. Using the reductions to monadic schemas, it suffices to show this for transformations that have monadic input schemas as input and output.

Fix a Δ_0 interpretation $\mathcal{I}$ with input $\mathfrak{U}_n$ and output $\mathfrak{U}_m$.

Before we proceed, first note that for every $d \leq m$, there is an NRC expression

$$E_d : \mathfrak{U}_n \to \mathsf{Set}(\mathfrak{U}_d)$$

collecting all of the subobjects of its input of sort $\mathfrak{U}_d$. It is formally defined by the induction over $n - d$.

$$E_m(x) \; := \; \{x\} \qquad\qquad E_d(x) = \bigcup E_{d-1}(x)$$

Write $E_{d_1,\dots,d_k}(x)$ for $\langle E_{d_1}, \dots, E_{d_k}\rangle(x)$ for every tuple of integers $d_1 \dots d_k$.

For $d \leq m$, let $d_1, \dots, d_k$ be the tuple such that the output sort $\mathfrak{U}_d$ is interpreted by the list of input sorts $\mathfrak{U}_{d_1}, \dots, \mathfrak{U}_{d_k}$. By induction over d, we build NRC expressions

$$E_d : \mathfrak{U}_m, \mathfrak{U}_{d_1}, \dots, \mathfrak{U}_{d_k} \to \mathfrak{U}_d$$

such that, provided that $\varphi_\mathsf{Domain}^{\mathfrak{U}_d}(\vec{a})$ and $\varphi_\mathsf{Domain}^{\mathfrak{U}_{d+1}}(\vec{b})$ hold, we have

$$\varphi_\in^{\mathfrak{U}_d}(\vec{a}, \vec{b}) \qquad\qquad \text{if and only if} \qquad\qquad E_d(\vec{a}) \in E_{d+1}(\vec{b})$$

For $E_0 : \mathfrak{U}_m, \mathfrak{U} \to \mathfrak{U}$, we simply take the second projection. Now assume that E_d is defined and that we are looking to define E_{d+1}. We want to set

$$E_{d+1}(x_{in}, \vec{y}) := \{E_d(x_{in}, \vec{x}) \mid \vec{x} \in E_{d_1,\dots,d_k}(x_{in}, \vec{y}) \wedge \mathsf{Verify}_{\varphi_\in^i}(x_{in}, \vec{x}, y_{in}, \vec{y})\}$$

which is NRC-definable as follows

$$\bigcup \Big\{ \mathsf{case}\left(\mathsf{Verify}_{\varphi_\in^i}(x_{in}, \vec{x}, y_{in}, \vec{y}), \; \{E_d(x_{in})\}, \; \{\} \right) \mid \vec{x} \in E_{d_1,\dots,d_k}(x_{in}) \Big\}$$

where Verify is given as in the Verification Proposition proven earlier in the supplementary materials and $\{E(\vec{x}, \vec{y}) \mid \vec{x} \in E'(\vec{y})\}$ is a notation for $\bigcup\{\ldots \bigcup\{E(\vec{x}, \vec{y}) \mid x_1 \in \pi_1(E'(\vec{y}))\} \ldots \mid x_k \in \pi_k(E'(\vec{y}))\}$. It is easy to check that the inductive invariant holds.

Now, consider the transformation $E_m : \mathfrak{U}_n, \mathfrak{U}_{m_1}, \ldots, \mathfrak{U}_{m_k} \to \mathfrak{U}_m$. The transformation

$$R \; := \; \{E_m(x_{in}, \vec{y}) \mid \vec{y} \in E_{m_1,\ldots,m_k}(x_{in}) \wedge \varphi^{\mathfrak{U}_m}_{\mathsf{Domain}}(\vec{y})\}$$

is also NRC-definable using Verify. Since the inductive invariant holds at level m, R returns the singleton containing the output of $\mathcal{I}$. Therefore $\mathsf{NRC}(R) : \mathfrak{U}_n \to \mathfrak{U}_m$ is the desired NRC expression equivalent to the interpretation $\mathcal{I}$. □

Note that the argument can be easily modified to produce an NRC expression that is *composition-free*: in union expressions $\bigcup\{E_1 \mid x \in E_2\}$, the range E_2 of the variable x is always another variable. In composition-free expressions, we allow as a native construct $\mathsf{case}(B, E_1, E_2)$ where B is a Boolean combination of atomic transformations with Boolean output, since we can not use composition to derive the conditional from the other operations.

Thus every NRC expression can be converted to one that is composition-free, and similarly for NRC. The analogous statements have been observed before for related languages like XQuery [BK09].

Appendix F. Completeness of proof systems

In the body of the paper we mentioned that the completeness of the high-level proof system of Figure 2 can be argued using a standard method – see, for example [Smu68a]. We outline this for the reader who is unfamiliar with these arguments. In Appendix G we will prove that the lower level system is complete, by reducing it to completeness of the higher level system.

We turn to the outline of the general method. One has a sequent $\Theta; \Gamma \vdash \Delta$ that is not provable. We want to construct a countermodel: one that satisfies all the formulas in Θ and Γ but none of the formulas in Δ. We construct a tree with $\Theta; \Gamma \vdash \Delta$ at the root by iteratively applying applicable inference rules in reverse: in "proof search mode", generating subgoals from goals. We apply the rules whenever possible in a given order, enforcing some fairness constraints: a rule that is active must be eventually applied along an infinite search, and if a choice of terms must be made (as with the ∃-R rule), all possible choices of terms are eventually made in an application of the rule. For example, if we have a disjunction $\rho_1 \vee \rho_2$ on the right, we may immediately "apply ∨-R": we generate a subgoal where on the right hand side we add ρ_1, ρ_2. Finite branches leading to a sequent that does not match the conclusion of any rule or axiom are artificially extended to infinite branches by repeating the topmost sequent.

By assumption, this process does not produce a proof, and thus we have an infinite branch b of the tree. We create a model M_b whose elements are the closure under tupling of the variables of types, other than product types, that appear on the branch. The model interprets these "non-tuple" variables by themselves, while variables with a product type are interpreted by – possibly nested – tuples of "non-tuple" variables. Specifically, M_b interprets terms t by these domain elements as follows.

- If t is a variable of type $\mathfrak{U}$ or of a set type, then t^{M_b} is t.
- If t is a variable of type Unit or is the term $\langle\rangle$, then t^{M_b} is $\langle\rangle$.

- If t is a variable of a product type, then t^{M_b} is $\langle x_1^{M_b}, x_2^{M_b} \rangle$, where t in the role of x and x_1, x_2 are the respective parameters of an instance of rule $\times_\eta$ on the branch. By our assumption that all applicable rules are applied in some fair way, there must be such an instance. Since it replaces all occurrences of t with a pair of fresh variables, there is exactly one such application, such that x_1 and x_2 are uniquely determined.
- If t is a term $\langle t_1, t_2 \rangle$, then $t^{M_b} = \langle t_1^{M_b}, t_2^{M_b} \rangle$.
- If t is a term $\pi_1(t')$, then t^{M_b} is x_1, where $t'^{M_b} = \langle x_1, x_2 \rangle$.
- If t is a term $\pi_2(t')$, then t^{M_b} is x_2, where $t'^{M_b} = \langle x_1, x_2 \rangle$.

The memberships correspond to the membership atoms that appear on the left of any sequent in b, and also the atoms that appear negated on the right hand side of any sequent.

We claim that M_b is the desired countermodel. It suffices to show that for every sequent $\Theta; \Gamma, \vdash \Delta$ in b, M_b is a counterexample to the sequent: it satisfies the conjunction of formulas on the left and none of the formulas on the right. We prove this by induction on the logical complexity of the formula. Each inductive step will involve the assumptions about inference rules not terminating proof search. For example, suppose for some sequent b_i in b of the above form, Δ contains $\rho_1 \vee \rho_2$, we want to show that M_b satisfies $\neg(\rho_1 \vee \rho_2)$.

But we know that in some successor (viewed root-first) b_j of b_i, we would have applied $\vee$-R, and thus have a descendant with ρ_1', ρ_2' within the right, where ρ_1' is identical to ρ_1 and ρ_2' to ρ_2, except possibly modified by applications of rules $\times_\eta$ and $\times_\beta$ in between b_i and b_j. By induction M_b satisfies $\neg\rho_1'$ and $\neg\rho_2'$. The rules $\times_\eta$ and $\times_\beta$ just perform replacement of terms, where the replaced and the replacing terms have the same denotation in M_b. Hence M_b also satisfies $\neg\rho_1$ and $\neg\rho_2$. Thus M_b satisfies $\neg(\rho_1 \vee \rho_2)$ as desired. The other connectives and quantifiers are handled similarly.

Appendix G. Completeness – translating to EL-normalized form

We show the completeness of our EL-normalized calculus of Figure 3 by giving a translation from proofs in the high-level calculus of Figure 2, whose completeness can be established with standard techniques, as outlined in Appendix F. We will give a precise account of the translation that starts from the following *base calculus*.

Definition G.1. Define the *base calculus* for Δ_0 formulas as the following modification of the EL-normalized calculus of Figure 3.

(1) Remove all EL decorations, i.e. do not require any contexts or sequents to be EL.
(2) In the $\exists$ rule drop the requirement that t is a tuple-term.

This base calculus may be seen as an intermediate between the high-level and the EL-normalized system. The high-level calculus in essence just provides some sugar on the base calculus, such that translation between both is straightforward. On the other hand, the base calculus is a relaxed version of the EL-normalized calculus: if a base calculus proof satisfies the respective EL requirements and all instances of the $\exists$ rule have as t a tuple-term, the proof can be considered as a proof in the EL-normalized calculus. The objective of our translation is then to enforce these constraints. The core technique for achieving the EL requirements is *permutation of rules*, investigated for standard systems in [Kle52] and [TS00, Section 5.3], adapted here to our proof system and applied in specific ways. Of particular relevance for us is the interaction of rule permutation with the following property of sequents in a proof.

Definition G.2. A top-level occurrence of a logic operator $\vee$, $\wedge$, or $\forall$ in a proof sequent is called *rule-dominated* if in all branches through the sequent it is a descendant[2] of the principal formula of an instance of the respective rule that introduces the operator. A proof is called $\vee\wedge\forall$-*rule-dominated* if the top-level occurrences of $\vee$, $\wedge$ and $\forall$ in all its sequents are rule dominated.

Recall that a multiset of formulas is EL if its members are positive atoms, negative atoms, existential quantifications and the truth-value constant $\bot$. In other words, it does not contain formulas with $\vee$, $\wedge$, $\forall$ or $\top$ as top-level operator. The following lemma is then easy to see.

Lemma G.3. *In an $\vee\wedge\forall$-rule-dominated proof, any sequent that is not below[3] an instance of one of the rules $\vee$, $\wedge$, or $\forall$ has the EL property.*

Correspondingly, our translation achieves the EL requirements by first extending the upper part of the given proof to bring it into $\vee\wedge\forall$-rule-dominated form and then permuting $\vee$, $\wedge$, or $\forall$ down over $\exists$, $\neq$, $\times_\eta$ and $\times_\beta$, where the $\vee\wedge\forall$-rule-dominated property is preserved. In a final step, parts of the proof are rewritten with special transformations to ensure that in all instances of the $\exists$ rule the term t is a tuple-term. Figure 4 specifies the steps of this method in detail.

To make claims about the involved rule permutations precise, we define *permutable* as follows.

Definition G.4. Rule R_α (with one or two premises) is said to be *permutable down over* rule R_β (with one premise) if for all instances α of R_α and β of R_β such that

(1) α is immediately above β where the conclusion of α is the premise of β,
(2) the principal formulas of α are disjoint with the active formulas of β,

α or β is void, i.e. has a single premise that is identical to its conclusion, or there is a deduction from the premises of α to the conclusion of β that consists of instances β_i' of R_β, one for each premise of α, whose conclusions are the premises of an instance α' of R_α, where α' has the same conclusion as β.

Also more general versions of *permutability* are possible, where, e.g., permutation of a one-premise rule down over a multi-premise rule like $\wedge$ is considered [TS00], but not required for our purposes.

Permutability properties of the rules of our base calculus, along with proof transformations to achieve the tuple-term property required for instances of the $\exists$ rule, justify the following claim.

Theorem G.5 (Completeness of the EL-normalized calculus for Δ_0 formulas (Figure 3)). *A Δ_0 formula provable in the base calculus for Δ_0 formulas (Def. G.1) is also provable in the EL-normalized calculus for Δ_0 formulas (Figure 3).*

Proof. Figure 4 shows a procedure for translating a proof in the base calculus to a proof in the EL-normalized calculus. Both proofs have the same bottom sequent, that is, prove the same Δ_0 formula. For step 1 of the procedure it is evident that the result is still a legal proof with the same bottom sequent and is in $\vee\wedge\forall$-rule-dominated form.

---

[2]In the non-strict sense: an occurrence is a descendant of itself.

[3]In the non-strict sense: the conclusion of a rule instance is below the rule instance.

Input: A proof in the base calculus for Δ_0 formulas (Def. G.1).

Method:

(1) *Bringing to $\vee\wedge\forall$-rule-dominated form by extending upwards at the top.* Bring the proof into $\vee\wedge\forall$-rule-dominated form by exhaustively extending its top nodes as follows: Select a top node (instance of $=$ or $\top$) whose sequent contains a formula occurrence with $\vee$, $\wedge$ or $\forall$ as top-level operator. Attach an instance of the rule for the respective operator such that the former top sequent becomes its conclusion, with the formula occurrence as principal formula.

(2) *Permuting rules to achieve the EL condition.* Permute all instances of $\vee$, $\wedge$ and $\forall$ down over instances of $\exists$, $\neq$, $\times_\eta$ and $\times_\beta$. Replace all derivations of sequents that have $\top$ as a member with applications of the $\top$ rule. Add the EL superscript to the contexts of the instances of $\exists$ and $\neq$ as well as to the sequents in instances of $\times_\eta$ and $\times_\beta$, as required by the EL-normalized calculus.

(3) *Bringing terms t in all instances of rule $\exists$ to tuple-term form.* As long as there is an occurrence of π_1 or π_2 in a term t of an instance of $\exists$, rewrite the proof part rooted at that instance with proof transformation *Elim-π-var* or *Elim-π-pair*.

Output: A proof in the EL-normalized calculus for Δ_0 formulas (Figure 3) with the same bottom sequent as the input proof.

Figure 4: EL-Normalization procedure.

The rule permutations in step 2 are justified by Lemma G.6 and Lemma G.7, presented below. It is easy to see from the particular deduction conversions used to justify these lemmas that the permutations preserve the $\vee\wedge\forall$-rule-dominated property of the proof. From Lemma G.3, the EL property required by the EL-normalized calculus for instances of $\exists$, $\neq$, $\times_\eta$ and $\times_\beta$ then follows. Step 2 finishes with straightforward conversions to let the proof literally match with the EL-normalized calculus: truncating redundant proof parts above sequents with $\top$ as a member and adding the EL decoration.

Step 3 is justified by Lemma G.8, presented below in Section G.3. The overall bottom sequent with the proven formula is unaltered by all involved permutations and transformations. The proof is then a proof in the EL-normalized calculus with the same bottom sequent as the initially given proof. $\square$

In the following Sects. G.1–G.3 we supplement the permutation lemmas and proof transformations referenced in Fig 4 and in the proof of Theorem G.5.

G.1. **Permutability justified by generic permutation schemas.** We begin our justification of the permuting of rules, used in the normalization algorithm. Permutability among rules that correspond to logic operators can be shown with straightforward generic schemas. The following lemmas express cases that were used in the proof of Theorem G.5.

Lemma G.6. *In the base calculus $\vee$, $\wedge$ and $\forall$ are permutable down over $\exists$ and $\neq$.*

Lemma G.6 hold since all of the involved permutations can be performed according to two generic schemas, one for permuting a one-premise rule down over a one-premise rule and a second schema for permuting the two-premise rule $\wedge$ down over a one-premise rule.

The considered one-premise rules $\neq, \vee, \forall, \exists$ are all of the form

$$\frac{\Theta, C \vdash A, \Delta}{\Theta, C' \vdash A', \Delta,} \tag{G.1}$$

where C, C' are (possibly empty) multisets of membership atoms and A, A' are multisets of formulas. For instances of the rule, C and A constitute the *active formulas* and C', A' the *principal formulas*.

Consider an instance α of a one-premise rule R_α followed by an instance β of a one-premise rule R_β, both rules of the form (G.1). Let the active formulas of the rule instances be C, A and D, B, respectively, and let their principal formulas be C', A' and D', B', respectively. Assume that the multiset A' of principal formulas of α and the multiset B of active formulas of β have no formula occurrences in common. (C' and D are not constrained in this way.) The two rule applications then match the left side of the following permutation schema. Except in the case where $R_\alpha = \exists$ and $R_\beta = \forall$, the schema can be applied to permute the instance of R_α down over that of R_β.

$$R_\alpha \frac{\Theta, C, D \vdash A, B, \Delta}{R_\beta \frac{\Theta, C', D \vdash A', B, \Delta}{\Theta, C', D' \vdash A', B', \Delta}} \quad \text{becomes} \quad R_\beta \frac{\Theta, C, D \vdash A, B, \Delta}{R_\alpha \frac{\Theta, C, D' \vdash A, B', \Delta}{\Theta, C', D' \vdash A', B', \Delta}} \tag{G.2}$$

Permutation schema (G.2) can for example be applied to permute $\forall$ down over $\exists$:

$$\exists \frac{\forall \frac{\Theta, t \in c, y \in b \vdash \varphi[y/x], \psi[t/z], \exists z \in c . \psi, \Delta}{\Theta, t \in c \vdash \forall x \in b . \varphi, \psi[t/z], \exists z \in c . \psi, \Delta}}{\Theta, t \in c \vdash \forall x \in b . \varphi, \exists z \in c . \psi, \Delta}$$

becomes

$$\forall \frac{\exists \frac{\Theta, t \in c, y \in b \vdash \varphi[y/x], \psi[t/z], \exists z \in c . \psi, \Delta}{\Theta, t \in c, y \in b \vdash \varphi[y/x], \exists z \in c . \psi, \Delta}}{\Theta, t \in c \vdash \forall x \in b . \varphi, \exists z \in c . \psi, \Delta}$$

The excluded case of permuting $\exists$ down over $\forall$ is not required for our purposes. To illustrate how this case might fail, consider the following instantiation, which shows two legal inferences matching the left side of permutation schema (G.2), but clearly not permitting permutation.

$$\forall \frac{\exists \frac{\Theta, y \in b \vdash \varphi[y/x], \psi[y/z], \exists z \in b . \psi, \Delta}{\Theta, y \in b \vdash \varphi[y/x], \exists z \in b . \psi, \Delta}}{\Theta \vdash \forall x \in b . \varphi, \exists z \in b . \psi, \Delta}$$

An instance of the two-premise rule $\wedge$ can be permuted below an instance of a one-premise rule R_β of the form (G.1) with active formulas D, B and principal formulas D', B' according to the following permutation schema.

$$R_\beta \frac{\wedge \frac{\Theta, D \vdash \varphi_1, B, \Delta \qquad \Theta, D \vdash \varphi_2, B, \Delta}{\Theta, D \vdash (\varphi_1 \wedge \varphi_2), B, \Delta}}{\Theta, D' \vdash (\varphi_1 \wedge \varphi_2), B', \Delta}$$

becomes   (G.3)

$$\land \frac{R_\beta \dfrac{\Theta, D \vdash \varphi_1, B, \Delta}{\Theta, D' \vdash \varphi_1, B', \Delta} \quad R_\beta \dfrac{\Theta, D \vdash \varphi_2, B, \Delta}{\Theta, D' \vdash \varphi_2, B', \Delta}}{\Theta, D' \vdash (\varphi_1 \land \varphi_2), B', \Delta}$$

As an example for the permutation schema (G.3) consider permuting $\land$ down over $\neq$:

$$\neq \frac{\land \dfrac{\Theta \vdash \varphi_1, t \neq_{\mathfrak{A}} u, \alpha[u/z], \alpha[t/z], \Delta \quad \Theta \vdash \varphi_2, t \neq_{\mathfrak{A}} u, \alpha[u/z], \alpha[t/z], \Delta}{\Theta \vdash \varphi_1 \land \varphi_2, t \neq_{\mathfrak{A}} u, \alpha[u/z], \alpha[t/z], \Delta}}{\Theta \vdash \varphi_1 \land \varphi_2, t \neq_{\mathfrak{A}} u, \alpha[t/z], \Delta}$$

becomes

$$\land \frac{\neq \dfrac{\Theta \vdash \varphi_1, t \neq_{\mathfrak{A}} u, \alpha[u/z], \alpha[t/z], \Delta}{\Theta \vdash \varphi_1, t \neq_{\mathfrak{A}} u, \alpha[t/z], \Delta} \quad \neq \dfrac{\Theta \vdash \varphi_2, t \neq_{\mathfrak{A}} u, \alpha[u/z], \alpha[t/z], \Delta}{\Theta \vdash \varphi_2, t \neq_{\mathfrak{A}} u, \alpha[t/z]\Delta}}{\Theta \vdash \varphi_1 \land \varphi_2, t \neq_{\mathfrak{A}} u, \alpha[t/z], \Delta}$$

G.2. **Permutability down over pairing and projection.** We now continue the justification of the rule permutations used in proof normalization. The rules $\times_\eta$ and $\times_\beta$ do not fit the generic shape (G.1) underlying the generic permutation schemas (G.2) and (G.3). Hence permuting a rule that corresponds to a logic operator down under $\times_\eta$ or $\times_\beta$ requires dedicated schemas. We consider rule $\times_\beta$ there just for π_1 (for π_2 the analogy is obvious). The following lemma expresses the cases of permutability over $\times_\eta$ and $\times_\beta$ needed in the proof of Theorem G.5.

Lemma G.7. *In the base calculus $\lor$, $\land$ and $\forall$ are permutable down over $\times_\eta$ and $\times_\beta$.*

The permutation schemas for $\lor$ and $\land$ down over $\times_\eta$ and $\times_\beta$ are straightforward. In their specification we use the following shorthands for a formula or multiset of formulas Γ: Γ^* stands for $\Gamma[\langle x_1, x_2 \rangle / x]$, Γ' for $\Gamma[t_1/x]$ and Γ'' for $\Gamma[\pi_1(\langle t_1, t_2 \rangle)/x]$. The respective permutation schemas then are as follows.

$$\times_\eta \frac{\lor \dfrac{\Theta^* \vdash \varphi_1^*, \varphi_2^*, \Delta^*}{\Theta^* \vdash \varphi_1^* \lor \varphi_2^*, \Delta^*}}{\Theta \vdash \varphi_1 \lor \varphi_2, \Delta} \quad \text{becomes} \quad \lor \frac{\times_\eta \dfrac{\Theta^* \vdash \varphi_1^*, \varphi_2^*, \Delta^*}{\Theta \vdash \varphi_1, \varphi_2, \Delta}}{\Theta \vdash \varphi_1 \lor \varphi_2, \Delta}$$

$$\times_\eta \frac{\land \dfrac{\Theta^* \vdash \varphi_1^*, \Delta^* \quad \Theta^* \vdash \varphi_2^*, \Delta^*}{\Theta^* \vdash \varphi_1^* \land \varphi_2^*, \Delta^*}}{\Theta \vdash \varphi_1 \land \varphi_2, \Delta} \quad \text{becomes} \quad \land \frac{\times_\eta \dfrac{\Theta^* \vdash \varphi_1^*, \Delta^*}{\Theta \vdash \varphi_1, \Delta} \quad \times_\eta \dfrac{\Theta^* \vdash \varphi_2^*, \Delta^*}{\Theta \vdash \varphi_2, \Delta}}{\Theta \vdash \varphi_1 \land \varphi_2, \Delta}$$

$$\times_\beta \frac{\lor \dfrac{\Theta' \vdash \varphi_1', \varphi_2', \Delta'}{\Theta' \vdash \varphi_1' \lor \varphi_2', \Delta'}}{\Theta'' \vdash \varphi_1'' \lor \varphi_2'', \Delta''} \quad \text{becomes} \quad \lor \frac{\times_\beta \dfrac{\Theta' \vdash \varphi_1', \varphi_2', \Delta'}{\Theta'' \vdash \varphi_1'', \varphi_2'', \Delta''}}{\Theta'' \vdash \varphi_1'' \lor \varphi_2'', \Delta''}$$

$$\times_\beta \frac{\land \dfrac{\Theta' \vdash \varphi_1', \Delta' \quad \Theta' \vdash \varphi_2', \Delta'}{\Theta' \vdash \varphi_1' \land \varphi_2', \Delta'}}{\Theta'' \vdash \varphi_1'' \land \varphi_2'', \Delta} \quad \text{becomes} \quad \land \frac{\times_\beta \dfrac{\Theta' \vdash \varphi_1', \Delta'}{\Theta'' \vdash \varphi_1'', \Delta''} \quad \times_\beta \dfrac{\Theta' \vdash \varphi_2', \Delta'}{\Theta'' \vdash \varphi_2'', \Delta''}}{\Theta'' \vdash \varphi_1'' \land \varphi_2'', \Delta''}$$

Permuting $\forall$ down over $\times_\eta$ and $\times_\beta$ is more intricate, because of the potential interaction of involved variables and substitutions. We make here explicit use of the assumption w.l.o.g. that bound and free variables are kept distinct.

For permuting $\forall$ down over $\times_\eta$ the given deduction is, under the assumption $x \neq z$, as follows.

$$\times_\eta \cfrac{\forall \cfrac{\Theta[\langle z_1, z_2 \rangle/z], y \in b[\langle z_1, z_2 \rangle/z] \vdash \varphi[\langle z_1, z_2 \rangle/z, y/x], \Delta[\langle z_1, z_2 \rangle/z]}{\Theta[\langle z_1, z_2 \rangle/z] \vdash \forall x \in b[\langle z_1, z_2 \rangle/z] \,.\, \varphi[\langle z_1, z_2 \rangle/z], \Delta[\langle z_1, z_2 \rangle/z]}}{\Theta \vdash \forall x \in b \,.\, \varphi, \Delta} \qquad \text{(G.4)}$$

We can exclude the cases $x = z_1$ or $x = z_2$, which would violate the freshness requirement of $\times_\eta$ because x already appears in the bottom sequent.

Assuming $x \neq z$, $x \neq z_1$ and $x \neq z_2$ it holds that

$$\varphi[\langle z_1, z_2 \rangle/z, y/x] = \varphi[y/x, \langle z_1, z_2 \rangle/z]$$

and thus (G.4) becomes

$$\forall \cfrac{\times_\eta \cfrac{\Theta[\langle z_1, z_2 \rangle/z], y \in b[\langle z_1, z_2 \rangle/z] \vdash \varphi[y/x, \langle z_1, z_2 \rangle/z], \Delta[\langle z_1, z_2 \rangle/z]}{\Theta, y \in b \vdash \varphi[y/x], \Delta}}{\Theta \vdash \forall x \in b \,.\, \varphi, \Delta} \qquad \text{(G.5)}$$

So far, for permuting $\forall$ down over $\times_\eta$ we assumed $x \neq z$. It remains to consider the case $x = z$. Since x occurs bound in the bottom sequent, it cannot occur in Θ, Δ and b. Since φ is in that sequent in the scope of a quantification upon x, the $\times_\eta$ inference cannot have any effect on φ. Hence, in the case $x = z$ the $\times_\eta$ inference would have no effect at all. Just deleting this void inference is then the result of permuting. This concludes the specification of permuting $\times_\eta$ down over $\forall$.

We now turn to permuting $\forall$ down over $\times_\beta$. Assuming w.l.o.g. that z does not occur in any sequent of the proof and thus $x \neq z$, the given deduction is as follows.

$$\times_\beta \cfrac{\forall \cfrac{\Theta[t_1/z], y \in b[t_1/z] \vdash \varphi[t_1/z, y/x], \Delta[t_1/z]}{\Theta[t_1/z] \vdash \forall x \in b[t_1/z] \,.\, \varphi[t_1/z], \Delta[t_1/z]}}{\Theta[\pi_1(\langle t_1, t_2 \rangle)/z] \vdash \forall x \in b[\pi_1(\langle t_1, t_2 \rangle)/z] \,.\, \varphi[\pi_1(\langle t_1, t_2 \rangle)/z], \Delta[\pi_1(\langle t_1, t_2 \rangle)/z]} \qquad \text{(G.6)}$$

In the case where x occurs neither in t_1 nor in t_2 it holds that

$$\varphi[t_1/z, y/x] = \varphi[y/x, t_1/z]$$
$$\varphi[y/x, \pi_1(\langle t_1, t_2 \rangle)/z] = \varphi[\pi_1(\langle t_1, t_2 \rangle)/z, y/x]$$

and thus (G.6) becomes

$$\forall \cfrac{\times_\beta \cfrac{\Theta[t_1/z], y \in b[t_1/z] \vdash \varphi[y/x, t_1/z], \Delta[t_1/z]}{\Theta[\pi_1(\langle t_1, t_2 \rangle)/z], y \in b[\pi_1(\langle t_1, t_2 \rangle)/z] \vdash \varphi[\pi_1(\langle t_1, t_2 \rangle)/z, y/x], \Delta[t_1/z]}}{\Theta[\pi_1(\langle t_1, t_2 \rangle)/z] \vdash \forall x \in b[\pi_1(\langle t_1, t_2 \rangle)/z] \,.\, \varphi[\pi_1(\langle t_1, t_2 \rangle)/z], \Delta[\pi_1(\langle t_1, t_2 \rangle)/z]}$$

We now consider the case where x occurs in t_1 or in t_2. Since x occurs bound in the bottom sequent, there can be no free occurrences of x in that sequent. Hence z, mapped in

that sequent to a pair in which x occurs, cannot occur in Θ, Δ and b. The given deduction (G.6) then specializes to

$$
\begin{array}{c}
\forall \dfrac{\Theta, y \in b \vdash \varphi[t_1/z, y/x], \Delta}{\Theta \vdash \forall x \in b \,.\, \varphi[t_1/z], \Delta} \\[2ex]
\times_\beta \dfrac{}{\Theta \vdash \forall x \in b \,.\, \varphi[\pi_1(\langle t_1, t_2\rangle)/z], \Delta}
\end{array}
\tag{G.7}
$$

Since $x \neq z$ it holds that

$$
\varphi[t_1/z, y/x] = \varphi[y/x, t_1[y/x]/z]
\tag{G.8}
$$
$$
\varphi[y/x, \pi_1(\langle t_1, t_2\rangle)[y/x]/z] = \varphi[\pi_1(\langle t_1, t_2\rangle)/z, y/x]
\tag{G.9}
$$

and thus (G.7) becomes

$$
\begin{array}{c}
\times_\beta \dfrac{\Theta, y \in b \vdash \varphi[y/x, t_1[y/x]/z], \Delta}{\Theta, y \in b \vdash \varphi[\pi_1(\langle t_1, t_2\rangle)/z, y/x], \Delta} \\[2ex]
\forall \dfrac{}{\Theta \vdash \forall x \in b \,.\, \varphi[\pi_1(\langle t_1, t_2\rangle)/z], \Delta}
\end{array}
\tag{G.10}
$$

In (G.10) the $\times_\beta$ inference viewed top-down replaces occurrences of $t_1[y/x]$, which, in case $t_1 = x$ is y, with $\pi_1(\langle t_1, t_2\rangle)[y/x]$. Justified by (G.9), the middle sequent is shown in (G.10) in the form suitable as premise for the application of $\forall$.

This concludes the specification of permuting $\forall$ down over $\times_\beta$ and thus the specification of all permutation schemas required for step 2 of the EL-normalization procedure (Figure 4).

G.3. **Conversion to tuple-terms.** The objective of this conversion is to ensure that the term t in instances of the $\exists$ rule is always a tuple-term, i.e., does not involve the projection functions π_1 and π_2. This is achieved with proof transformations where instances of $\times_\eta$ and $\times_\beta$ are inserted below instances of $\exists$ and the effects of these interspersed rules are propagated upwards. The following lemma states the desired overall property of this transformation.

Lemma G.8. *A Δ_0 formula provable in the base calculus with a proof where the* EL *requirements of the* EL-*normalized calculus of Figure 3 are met (but not necessarily its tuple-term condition of the $\exists$ rule) is also provable in the* EL-*normalized calculus.*

Lemma G.8 presupposes a proof whose contexts and sequents meet the EL requirements of the EL-normalized calculus. The claimed tuple-term condition can be achieved by exhaustively applying the proof transformations *Elim-π-var* and *Elim-π-pair*, specified below, which are applicable whenever there is an occurrence of π_1 or π_2 in the term t of an instance of the $\exists$ rule. These transformations lead from a legal proof in the base calculus to another one with the same bottom formula. They require and preserve the EL requirements from the EL-normalized calculus.

Proof transformation Elim-π-var. This transformation applies to an instance of the $\exists$ rule whose term t has an occurrence of $\pi_1(z)$ or $\pi_2(z)$, where z is a variable. The instance of the $\exists$ rule is replaced with an instance of $\exists$ followed by an instance of $\times_\eta$ and the subproof

above is adjusted, according to the following schema.

$$\exists \; \frac{\Theta, t \in b \vdash \varphi[t/x], \exists x \in b \; \varphi, \Delta_{\mathsf{EL}}}{\Theta, t \in b \vdash \exists x \in b \, . \, \varphi, \Delta_{\mathsf{EL}}}$$

with subproof P above.

is replaced by

$$\times_\eta \; \frac{\exists \; \dfrac{\Theta^*, t^* \in b^* \vdash \varphi^*[t^*/x], \exists x \in b^* \; \varphi, \Delta^*_{\mathsf{EL}}}{\Theta^*, t^* \in b^* \vdash \exists x \in b^* \, . \, \varphi^*, \Delta^*_{\mathsf{EL}}}}{\Theta, t \in b \vdash \exists x \in b \, . \, \varphi, \Delta_{\mathsf{EL}}}$$

with subproof P^* above.

The starred components are specified there as follows: For a term, formula, or multiset of formulas Γ, Γ^* stands for $\Gamma[\langle z_1, z_2 \rangle / z]$, where z_1, z_2 are fresh variables. Subproof P^* is obtained from P by removing each instance of $\times_\eta$ upon the variable z and some pair $\langle z_1', z_2' \rangle$, while replacing and z_1' with z_1 and of z_2' with z_2, and then replacing all occurrences of z with $\langle z_1, z_2 \rangle$.

Proof transformation Elim-π-pair. This transformation applies to an instance of the $\exists$ rule whose term t has an occurrence of $\pi_1(\langle t_1, t_2 \rangle)$ (the case for $\pi_2(\langle t_1, t_2 \rangle)$ is analogous). The instance of the $\exists$ rule is replaced with an instance of $\exists$ followed by an instance of $\times_\beta$ and the subproof above is adjusted, according to the following schema.

$$\exists \; \frac{\Theta, t \in b \vdash \varphi[t/x], \exists x \in b \; \varphi, \Delta_{\mathsf{EL}}}{\Theta, t \in b \vdash \exists x \in b \, . \, \varphi, \Delta_{\mathsf{EL}}}$$

with subproof P above.

is replaced by

$$\times_\beta \; \frac{\exists \; \dfrac{\Theta, t' \in b \vdash \varphi[t'/x], \exists x \in b \; \varphi, \Delta_{\mathsf{EL}}}{\Theta, t' \in b \vdash \exists x \in b \, . \, \varphi, \Delta_{\mathsf{EL}}}}{\Theta, t \in b \vdash \exists x \in b \, . \, \varphi, \Delta_{\mathsf{EL}}}$$

with subproof P' above.

The primed components are specified there as follows: Term t' is t with the occurrence of $\pi_1(\langle t_1, t_2 \rangle)$ replaced by t_1. Subproof P' is P with that replacement propagated upward. It may be obtained by performing the inference steps of P in upward direction, but now starting from $\Theta, t' \in b \vdash \varphi[t'/x], \exists x \in b \; \varphi, \Delta_{\mathsf{EL}}$ instead of $\Theta, t \in b \vdash \varphi[t/x], \exists x \in b \; \varphi, \Delta_{\mathsf{EL}}$. Instances of $\times_\beta$ and $\times_\eta$ may get void there (identical premise and conclusion) and can then be removed.

G.4. **Complexity considerations.** In step 1 of the EL-normalization procedure (Figure 4), through the two-premise rule $\wedge$ a leaf can get extended upwards by a number of nodes that is exponential in the number of occurrences of $\wedge$ in its conclusion. Hence the time required for step 1 is exponential in the number of occurrences of $\wedge$ in a leaf conclusion.

Step 2 where the proof is permuted, may, through permuting $\wedge$ down over one-premise rules, increase the number of nodes. However, none of the permutations increases the height h or the number of leaves l of the proof tree. Both h and l are less than or equal to the number n of nodes of the input proof, the output of step 1. Since a tree with height h and number of leaves l can in general not have more than $h \times l$ nodes, our proof tree can never have more than n^2 nodes when we perform the permutations. Since for each node we can apply at most h downward permutation steps, and $h \leq n$, we can conclude that step 2 involves not more that n^3 permutation steps.

Step 3, where instances of $\times_\eta$ and $\times_\beta$ are inserted below instances of $\exists$ and adjustments to these insertions are propagated upwards, is polynomial. This step does not increase

the number of branches, but may increase the length of a branch by twice the number of occurrences of π_i in terms t of instances of $\exists$ on the branch.

Thus, for EL-normalization we have exponential time complexity. However, the source of the exponential complexity can be associated with the first stage of the procedure, extending the proof tree at its leaves, which may be seen as a blow up of *irrelevant* parts of the proof that accompany the axioms. The rule permutation stage is only a polynomial-time operation.

APPENDIX H. Δ_0 INTERPOLATION: PROOF SKETCH OF THEOREM 5.4

We recall the statement:

Let Θ be an $\in$-context and Γ, Δ finite multisets of Δ_0 formulas. Then from any proof of $\Theta;\ \Gamma \vdash \Delta$ we can compute in linear time a Δ_0 formula θ with $FV(\theta) \subseteq FV(\Theta,\Gamma) \cap FV(\Delta)$, such that $\Theta;\ \Gamma \vdash \theta$ and $\emptyset;\ \theta \vdash \Delta$

Recall also that we claim this for both the higher-level 2-sided system and the 1-sided system, where the 2-sided syntax is a "macro": $\Theta;\ \Gamma \vdash \Delta$ is a shorthand for $\Theta \vdash \neg\Gamma, \Delta$, where $\neg\Gamma$ is itself a macro for dualizing connectives. Thus in the 1-sided version, we are arbitrarily classifying some of the Δ_0 formulas as Left and the others as Right, and our interpolant must be common according to that partition.

We stress that there are no new ideas needed in proving Theorem 5.4 — unlike for our main tool, the Parameter Collection Theorem, or our final result. The construction for Theorem 5.4 proceeds exactly as in prior interpolation theorems for similar calculi [Tak87, TS00, Smu68b]. Similar constructions are utilized in works for query reformulation in databases, so for a presentation geared towards a database audience one can check [TW11] or the later [BCLT16].

We explain the argument for the higher-level 2-sided system. We prove a more general statement, where we partition the context and the formulas on both sides of $\vdash$ into Left and Right. So we have

$$\Theta_L\Theta_R; \Gamma_L, \Gamma_R \vdash \Delta_L, \Delta_R$$

And our inductive invariant is that we will compute in linear time a θ such that:

$$\Theta_L; \Gamma_L \vdash \theta, \Delta_L$$
$$\Theta_R; \Gamma_R, \theta \vdash \Delta_R$$

And we require that $FV(\theta) \subseteq FV(\Theta_L, \Gamma_L, \Delta_L) \cap FV(\Theta_R, \Gamma_R, \Delta_R)$. This generalization is used to handle the negation rules, as we explain below.

We proceed by induction on the depth of the proof tree.

One of the base cases is where we have a trivial proof tree, which uses rule (Ax) to derive:

$$\Theta; \Gamma, \varphi \vdash \varphi, \Delta$$

We do a case distinction on where the occurrences of φ sit in our partition. Assume the occurrence on the left is in Γ_L and the occurrence on the right is in Δ_R. Then we can take our interpolant θ to be φ. Suppose the occurrence on the left is Γ_L and the occurrence on the right is in Δ_L. Then we can take θ to be $\bot$. The other base cases are similar.

The inductive cases for forming the interpolant will work "in reverse" for each proof rule. That is, if we used an inference rule to derive sequent S from sequents S_1 and S_2, we will partition the sequents S_1 and S_2 based on the partition of S. We will then apply induction

to our partitioned sequent for S_1 to get an interpolant θ_1, and also apply induction to our partitioned version of S_2 to get an interpolant θ_2. We then put them together to get the interpolant for the partitioned sequent S. This "putting together" will usually reflect the semantics of the connective mentioned in the proof rule.

Consider the case where the last rule applied is the $\neg$-L rule: this is the case that motivates the more general invariant involving partitions. We have a partition of the final sequent $\Theta; \Gamma, \varphi \vdash \Delta$. We form a partition of the sequent $\Theta; \Gamma \vdash \neg\varphi, \Delta$ by placing $\neg\varphi$ on the same side (Left, Right) as φ was in the original partition. We then get an interpolant θ by induction. We just use θ for the final interpolant.

We consider the inductive case for $\wedge$-R. We have two top sequents, one for each conjunct. We partition them in the obvious way: each φ_i in the top is in the same partition that $\varphi_1 \wedge \varphi_2$ was in the bottom. Inductively we take the interpolants θ_1 and θ_2 for each sequent. We again do a case analysis based on whether $\varphi_1 \wedge \varphi_2$ was in Δ_L or in Δ_R.

Suppose $\varphi_1 \wedge \varphi_2$ was in Δ_R, so $\Delta_R = \varphi_1 \wedge \varphi_2, \Delta'_R$. Then we arranged that each φ_i was in Δ_R in the corresponding top sequent. So we know that $\Theta_L; \Gamma_L \vdash \theta_i, \Delta_L$ and $\Theta_R; \Gamma_R, \theta_i \vdash \varphi_i, \Delta'_R$ for $i = 1, 2$. Now we can set the interpolant θ to be $\theta_1 \wedge \theta_2$.

In the other case, $\varphi_1 \wedge \varphi_2$ was in Δ_L, say $\Delta_L = \varphi_1 \wedge \varphi_2, \Delta'_L$. Then we would arrange each φ_i to be "Left" in the corresponding top sequent, so we know that $\Theta_L; \Gamma_L \vdash \theta_i, \varphi_i, \Delta'_L$ and $\Theta_R; \Gamma_R, \theta_i \vdash \Delta_R$ for $i = 1, 2$. We set $\theta = \theta_1 \vee \theta_2$ in this case.

With the $\exists$ rule, a term in the inductively-assumed θ' for the top sequent may become illegal for the θ for the bottom sequent, since it has a free variable that is not common. In this case, the term in θ is replaced by a quantified variable, where the quantifier is existential or universal, depending on the partitioning, and bounded according to the requirements for Δ_0 formulas. The non-common variables are then replaced and variable constraints are added as described with the notation $\exists x_1 \ldots x_n | t \in b \,.\, \varphi$ and $\forall x_1 \ldots x_n | t \in b \,.\, \varphi$ on p. 30.

Appendix I. Proofs of polytime admissibility

The goal of this section is to prove most claims of polytime admissibility made in the body of the paper, crucially those of Subsection 5.3. Recall that a rule

$$\frac{\Theta \vdash \Delta}{\Theta' \vdash \Delta'}$$

is *polytime admissible* if we can compute in polynomial time a proof of the conclusion $\Theta' \vdash \Delta'$ from a proof of the antecedent $\Theta \vdash \Delta$.

Throughout this section we deal with the EL-normalized proof system of Figure 3.

I.1. **Standard rules.** Here we collect some useful standard sequent calculi rules, which are all polytime admissible in our system. The arguments for these rules are straightforward.

Lemma I.1. *The following weakening rule is polytime admissible:*

$$\frac{\Theta \vdash \Delta}{\Theta' \vdash \Delta, \Delta'}$$

Lemma I.2. *The following inference, witnessing the invertibility of the $\wedge$ rule, is polytime admissible for both $i \in \{1, 2\}$:*

$$\frac{\Theta \vdash \varphi_1 \wedge \varphi_2, \Delta}{\Theta \vdash \varphi_i, \Delta}$$

Lemma I.3. *The following, witnessing the invertibility of the $\forall$ rule, is polytime admissible:*
$$\frac{\Theta \vdash \forall x \in t.\varphi, \Delta}{\Theta, x \in t \vdash \varphi, \Delta}$$

Lemma I.4. *The following substitution rule is polytime admissible:*
$$\frac{\Theta \vdash \Delta}{\Theta[t/x] \vdash \Delta[t/x]}$$

I.2. **Admissibility of generalized congruence.** Recall the admissibility claim concerning the rule related to congruence:

Lemma I.5. *The following generalized congruence rule is polytime admissible:*
$$\frac{\Theta[t/x, t/y] \vdash \Delta[t/x, t/y]}{\Theta[t/x, u/y] \vdash \neg(t \equiv u), \Delta[t/x, u/y]}$$

Recall that in a two-sided reading of this, the hypothesis is $t \equiv u; \Theta[t/x, u/y] \vdash \Delta[t/x, u/y]$. So the rule says that if we Θ entails Δ where both contain t, then if we assume $t \equiv u$ and substitute some occurrences of t with u in Θ, we can conclude the corresponding substitution of Δ.

To prove Lemma I.5 in the case where the terms t and u are of type $\mathsf{Set}(T)$, we will need a more general statement. We are going to generalize the statement to treat tuples of terms and use EL formulas instead of $\neg(t \equiv u)$ to simplify the inductive invariant.

Given two terms t and u of type t, define by induction the set of formulas $\mathcal{E}_{t,u}$:

- If $T = \mathfrak{U}$, then $\mathcal{E}_{t,u}$ is $t \neq_{\mathfrak{U}} u$
- If $T = T_1 \times T_2$, then $\mathcal{E}_{t,u}$ is $\mathcal{E}_{\pi_1(t), \pi_1(u)}, \mathcal{E}_{\pi_2(t), \pi_2(u)}$
- If $T = \mathsf{Set}(T')$, $\mathcal{E}_{t,u}$ is $\neg(t \subseteq_T u), \neg(u \subseteq_T t)$

The reader can check that $\mathcal{E}_{t,u}$ is essentially $\neg(t \equiv u)$.

Lemma I.6. *The following rule is polytime admissible:*
$$\frac{\Theta \vdash \Delta, \mathcal{E}_{t,u}}{\Theta \vdash \Delta, \neg(t \equiv u)}$$

Proof. Straightforward induction over T. $\qquad\square$

Since we will deal with multiple equivalences, we will adopt vector notation $\vec{t} = t_1, \ldots, t_n$ and $\vec{x} = x_1, \ldots, x_n$ for lists of terms and variables. Call $\mathcal{E}_{\vec{t}, \vec{u}}$ the union of the $\mathcal{E}_{t_i, u_i}$. We can now state our more general lemma:

Lemma I.7. *The following generalized n-ary congruence rule for set variables is polytime admissible:*
$$\frac{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}]}{\Theta[\vec{t}/\vec{x}, \vec{u}/\vec{y}] \vdash \Delta[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \mathcal{E}_{\vec{t}, \vec{u}}}$$

Proof of Lemma I.7. We proceed by induction over the proof of $\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash \Delta[\vec{t}/\vec{x}, \vec{u}/\vec{y}]$,

- If the last rule applied is the = rule, i.e. we have

$$\overline{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash a[\vec{t}/\vec{x}, \vec{t}/\vec{y}] =_{\mathfrak{U}} b[\vec{t}/\vec{x}, \vec{t}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}]}$$

and $a[\vec{t}/\vec{x}, \vec{t}/\vec{y}] = b[\vec{t}/\vec{x}, \vec{t}/\vec{y}]$, with a, b variables. Now if a and b are equal, or if they belong both to either $\vec{x}$ or $\vec{y}$, it is easy to derive the desired conclusion with a single application of

the $=$ rule. Otherwise, assume $a = x_i$ and $b = y_j$ (the symmetric case is handled similarly). In such a case, we have that $\mathcal{E}_{\vec{t},\vec{u}}$ contains $t_i \neq_{\mathfrak{U}} t_j$. So the desired proof

$$\overline{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash t_i =_{\mathfrak{U}} u_j, \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}, t_i \neq_{\mathfrak{U}} t_j}$$

follows from the polytime admissibility of the axiom rule.

- Suppose the last rule applied is the $\top$ rule:

$$\overline{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash \top, \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}]}$$

Then we do not need to apply the induction hypothesis. Instead we can immediately apply the $\top$ rule to obtain

$$\overline{\Theta[\vec{t}/\vec{x}, \vec{u}/\vec{y}] \vdash \top, \Delta[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}}$$

- If the last rule applied is the $\wedge$ rule

$$\frac{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash \varphi_1[\vec{t}/\vec{x}, \vec{t}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \qquad \Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash \varphi_2[\vec{t}/\vec{x}, \vec{t}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}]}{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash (\varphi_1 \wedge \varphi_2)[\vec{t}/\vec{x}, \vec{t}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}]}$$

then the induction hypothesis gives us proofs of

$$\Theta[\vec{t}/\vec{x}, \vec{u}/\vec{y}] \vdash \varphi_i[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}$$

for both $i \in \{1, 2\}$. So we can apply the $\wedge$ rule to conclude that we have

$$\Theta[\vec{t}/\vec{x}, \vec{u}/\vec{y}] \vdash (\varphi_1 \wedge \varphi_2)[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}$$

as desired.

- The cases of the rules $\vee, \forall$ and $\times_\eta$ are equally straightforward and left to the reader.
- Now, let us handle the case of the $\exists$ rule.

$$\frac{\Theta, t \in u \vdash \varphi[t/x], \Delta^{\mathsf{EL}}}{\Theta, t \in u \vdash \exists x \in u.\ \varphi, \Delta^{\mathsf{EL}}}$$

So assume that z is fresh wrt $\vec{x}, \vec{y}, \vec{t}, \vec{u}, a, b$ and that the last step of the proof is

$$\frac{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}], a[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \in b[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash \varphi[a/z][\vec{t}/\vec{x}, \vec{t}/\vec{y}], \exists z \in c[\vec{t}/\vec{x}, \vec{t}/\vec{y}].\ \varphi[\vec{t}/\vec{x}, \vec{t}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}]}{\Theta[\vec{t}/\vec{x}, \vec{t}/\vec{y}], a[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \in b[\vec{t}/\vec{x}, \vec{t}/\vec{y}] \vdash \exists z \in c[\vec{t}/\vec{x}, \vec{t}/\vec{y}].\ \varphi[\vec{t}/\vec{x}, \vec{t}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{t}/\vec{y}]}$$

with $b[\vec{t}/\vec{x}, \vec{t}/\vec{y}] = c[\vec{t}/\vec{x}, \vec{t}/\vec{y}]$. Set $a' = a[\vec{t}/\vec{x}, \vec{u}/\vec{y}]$, $b' = b[\vec{t}/\vec{x}, \vec{u}/\vec{y}]$, $c' = c[\vec{t}/\vec{x}, \vec{u}/\vec{y}]$. We have three subcases:

- If we have that $b' = c'$, using the induction hypothesis, we have a proof of

$$\Theta[\vec{t}/\vec{x}, \vec{u}/\vec{y}], a' \in b' \vdash \varphi[a/z][\vec{t}/\vec{x}, \vec{u}/\vec{y}], \exists z \in c'.\ \varphi[\vec{u}/\vec{x}, \vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}$$

we can simply apply an $\exists$ rule to that proof and we are done.

- Otherwise, if we have that $b' = t_i$ and $c' = u_j$ for some $i, j \leq n$. In that case, extending the tuples $\vec{t}$ and $\vec{u}$ with a' and a fresh variable z' (the substitutions under consideration would be, we can apply the induction hypothesis to obtain a proof of

$$\Theta[\vec{t}/\vec{x}, \vec{u}/\vec{y}, a'/z], a' \in t_i, z' \in u_j \vdash$$
$$\varphi[\vec{t}/\vec{x}, \vec{u}/\vec{y}, z'/z], \exists z \in u_j.\ \varphi[\vec{u}/\vec{x}, \vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x}, \vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}, \mathcal{E}_{a',z'}$$

Note that $\mathcal{E}_{\vec{t},\vec{u}}$ contains an occurrence of $\neg(t_i \subseteq u_j)$, which expands to $\exists z \in t_i.\forall z' \in u_j.\neg(z \equiv z')$, so we can construct the partial derivation

$$
\begin{array}{c}
\cfrac{
\cfrac{
\cfrac{
\cfrac{
\Theta[\vec{t}/\vec{x},\vec{u}/\vec{y},a'/z], a' \in t_i, z' \in u_j \vdash \varphi[\vec{t}/\vec{x},\vec{u}/\vec{y},z'/z'], \exists z \in u_j.\ \varphi[\vec{u}/\vec{x},\vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x},\vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}, \mathcal{E}_{a',z'}
}{
\Theta[\vec{t}/\vec{x},\vec{u}/\vec{y}], a' \in t_i, z' \in u_j \vdash \exists z \in u_j.\varphi[\vec{t}/\vec{x},\vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x},\vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}, \mathcal{E}_{a',z'}
}\ \exists
}{
\Theta[\vec{t}/\vec{x},\vec{u}/\vec{y},a'/z], a' \in t_i, z' \in u_j \vdash \exists z \in u_j.\ \varphi[\vec{u}/\vec{x},\vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x},\vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}, a' \equiv z'
}\ \text{Lemma I.6}
}{
\Theta[\vec{t}/\vec{x},\vec{u}/\vec{y},a'/z], a' \in t_i \vdash \exists z \in c'.\ \varphi[\vec{u}/\vec{x},\vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x},\vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}, \forall z' \in u_j.a' \equiv z'
}\ \forall
}{
\Theta[\vec{t}/\vec{x},\vec{u}/\vec{y},a'/z], a' \in t_i \vdash \exists z \in c'.\ \varphi[\vec{t}/\vec{x},\vec{u}/\vec{y}], \Delta[\vec{t}/\vec{x},\vec{u}/\vec{y}], \mathcal{E}_{\vec{t},\vec{u}}
}\ \exists
\end{array}
$$

whose conclusion matches what we want.
- Otherwise, we are in a similar case where $b' = u_j$ and $c' = t_i$. We proceed similarly, except that we use the formula $\neg(u_j \subseteq t_i)$ of $\mathcal{E}_{t_i,u_j}$ instead of $\neg(t_i \subseteq u_j)$.
- For the rule $\times_\beta$, which has general shape

$$
\frac{\Theta[z_i/z] \vdash \Delta[z_i/z]}{\Theta[\pi_i(\langle z_1, z_2\rangle)/z] \vdash \Delta[\pi_i(\langle z_1, z_2\rangle)/z]}
$$

we can assume, without loss of generality, that z occurs only once in Θ, Δ. Let us only sketch the case where z occurs in a formula φ and the rule has shape

$$
\frac{\Theta \vdash \varphi[z_i/z], \Delta}{\Theta \vdash \varphi[\pi_i(\langle z_1, z_2\rangle)/z], \Delta}
$$

We have that $\varphi[\pi_i(\langle z_1, z_2\rangle)/z)]$ is also of the shape $\psi[\vec{t}/\vec{x}, \vec{u}/\vec{y}]$ in our situation. We can also assume without loss of generality that each variable in $\vec{x}$ and $\vec{y}$ occur each a single time in Θ, Δ. Now if we have a couple of situations:
* If the occurrence of z do not interfere with the substitution $[\vec{t}/\vec{x}, \vec{u}/\vec{y}]$, i.e., there is a formula θ such that

$$
\varphi[\pi_i(\langle z_1, z_2\rangle)/z] = \psi[\vec{t}/\vec{x}, \vec{u}] = \theta[\vec{t}/\vec{x}, \vec{u}/\vec{y}, \langle z_1, z_2\rangle/z]
$$

  we can simply apply the induction hypothesis on the subproof and conclude with one application of $\times_\eta$.
* If we have that z clashes with a variable of $\vec{x}, \vec{y}$, say x_j, but that $t_j = v[\pi_i(\langle z_1, z_2\rangle/x_j)]$ for some term v. Then we can apply the induction hypothesis with the matching tuples of terms $\vec{t}, x_i$ and $\vec{u}, t_i$ and conclude by applying the β rule.
* Otherwise, the occurrence of z does interfere with the substitution in such a way that we have, say $t_j = \langle z_1, z_2\rangle$. Then we can apply the induction hypothesis on the subproof with the matching tuples of terms $\vec{x}, z_i$ and $\vec{y}, \pi_i(u_j)$ and conclude by applying the β rule. $\qquad\square$

One easy consequence of the above is Lemma I.5:

Proof of Lemma I.5. Combine Lemma I.7 and Lemma I.6. $\qquad\square$

Another consequence is the following corollary, which will be used later in this section:

Corollary I.8. *The following rule is polytime admissible:*

$$
\frac{\Theta, t \in u \vdash \Delta}{\Theta \vdash \neg t \,\hat{\in}\, u, \Delta}
$$

Proof. Recall that $t \,\hat{\in}\, u$ expands to $\exists x \in u.\ x \equiv t$, so that $\neg t \,\hat{\in}\, u$ is $\forall x \in u.\ \neg(x \equiv t)$. So we have

$$\text{Lemma I.5} \; \frac{\dfrac{\Theta, t \in u \vdash \Delta}{\Theta, x \in u \vdash \neg(x \equiv t), \Delta}}{\forall \; \overline{\Theta \vdash \neg t \,\hat{\in}\, u, \Delta}}$$

$\square$

I.3. **Proof of Lemma 5.7.** We now recall the claim of admissibility concerning rules for "moving down in an equivalence". These will make use of the notation for quantifying on a path below an object, defined in the body in Definition 5.6.

Lemma 5.7. *Assume p is a subtype occurrence for the type of the term o'. The following is polytime admissible*

$$\frac{\Theta \vdash \Delta, \; \exists r' \in_p o'. \; r \equiv_{\mathsf{Set}(T')} r'}{\Theta, z \in r \vdash \Delta, \; \exists z' \in_{\ni,p} o'. \; z \equiv_{T'} z'}$$

Furthermore, the size of the output proof is at most the size of the input proof.

Proof. First, let us consider the simpler case where $p = \ni$. We proceed by induction over the input proof of $\Theta \vdash \Delta, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'$ and make a case distinction according to which rule was applied last. All cases are straightforward, save for one: when a $\exists$ rule is applied on the formula $\exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'$. Let us only detail that one.

In that case, the last step has shape

$$\frac{\Theta, w \in o' \vdash \Delta^{\mathsf{EL}}, r \equiv_{\mathsf{Set}(T')} w, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'}{\Theta, w \in o' \vdash \Delta^{\mathsf{EL}}, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'}$$

Now observe that $r \equiv_{\mathsf{Set}(T')} w$ is not EL, since if you unfold the macros it begins with a universal. Thus the final rule that is applied in the proof witnessing the hypothesis sequent cannot be the $\exists$ rule. We can thus infer that the expanded proof tree ends with:

$$\frac{\dfrac{\Theta, w \in o', z \in r \vdash \Delta^{\mathsf{EL}}, z \,\hat{\in}\, w, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'}{\Theta, w \in o' \vdash \Delta^{\mathsf{EL}}, r \subseteq w, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'} \quad \dfrac{\vdots}{\Theta, w \in o' \vdash \Delta^{\mathsf{EL}}, w \subseteq r, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'}}{\dfrac{\Theta, w \in o' \vdash \Delta^{\mathsf{EL}}, r \equiv_{\mathsf{Set}(T')} w, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'}{\Theta, w \in o' \vdash \Delta^{\mathsf{EL}}, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'}}$$

In particular, we have a strictly smaller subproof of

$$\Theta, w \in o', z \in r \vdash \Delta^{\mathsf{EL}}, z \,\hat{\in}\, w, \exists r' \in o'. \; r \equiv_{\mathsf{Set}(T')} r'$$

Recall that $\hat{\in}$ is a macro: see Definition 3.3. Applying the induction hypothesis, we get a proof of

$$\Theta, w \in o', z \in r \vdash \Delta^{\mathsf{EL}}, z \,\hat{\in}\, w, \exists z' \in_{\ni,\ni} o'. \; z \equiv_{\mathsf{Set}(T')} z'$$

Applying the $\exists$-rule we can then obtain a proof with conclusion

$$\Theta, w \in o', z \in r \vdash \Delta^{\mathsf{EL}}, \exists r \in o'. \; z \,\hat{\in}\, r, \exists z' \in_{\ni,\ni} o'. \; z \equiv_{\mathsf{Set}(T')} z'$$

which concludes our argument, since $\exists r \in_p o'. \; z \,\hat{\in}\, r$ and $\exists z' \in_{\ni,\ni} o'. \; z \equiv_{\mathsf{Set}(T')} z'$ are syntactically the same.

For the more complex case where p is non-trivial, we can conduct a similar argument, at the cost of making the induction hypothesis more elaborate. We would prove that the following rule is polytime admissible

$$\frac{\Theta, \Theta_1, \dots, \Theta_n \vdash \Delta, \varphi_1, \dots, \varphi_n}{\Theta, \Theta_1, \dots, \Theta_n \vdash \Delta, \exists z' \in_{\ni,p} o'. \; z \equiv_{\mathsf{Set}(T')} z'}$$

Here n is any natural number and for each $i \le n$, we have a decomposition of p as a concatenation p_i, p'_i, where p'_i is non-empty and multiple p_i can be the same. From this decomposition we define $\varphi_i = \exists z' \in_{p'_i} o_i.\ z \equiv_{\mathsf{Set}(T')} z'$ (with $o_i = o'$ iff $p_i = \varepsilon$) and Θ_i is a context witnessing that $o_i \in_{p_i} o'$. Formally, $\Theta_i = \Theta(p_i, o_i, o')$ with $\Theta(\varepsilon, x, y) = \cdot$, $\Theta(\ni, x, y) = x \in y$, $\Theta((p, \ni), x, y) = \Theta(p, x, z), z \in y$ and $\Theta((p, j), x, y) = \Theta(p, x, \pi_j(y))$ when $j \in \{1, 2\}$. The induction over a proof of the premise can be carried out in an analogous way to show admissibility, and the only non-trivial case, when a rule interacts with one of the φ_i that has $p'_i = \ni$, is handled in the same way as the non-trivial case for $p = \ni$. □

I.4. **Proof of Lemma 5.8.**

Lemma 5.8. *The following is polytime admissible (where p is a subtype occurrence of the type of o')*

$$\frac{\Theta \vdash \Delta, \exists r' \in_p o'.\ r \equiv_{\mathsf{Set}(T')} r'}{\Theta \vdash \Delta, \exists r' \in_p o'.\forall z \in a.\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'}$$

Proof. Much like what happened in the proof of 5.7, the proof can be done by induction over the shape of the input derivation when $p = \ni$. When this is not the case, we can generalize our inductive hypothesis in a similar way and use the same ideas. So for the sake of legibility, let us focus on the case where $p = \ni$ and go through the induction, making a case distinction according to what was the last rule applied.

All cases are trivial, except for the case of the rule $\exists$, where $\exists r' \in o'.r \equiv_{\mathsf{Set}(T')} r'$ is the principal formula. So let us focus on that one.

In that case, the proof necessarily has shape

$$\wedge\frac{\forall\dfrac{\Theta, y \in w \vdash \Delta^{\mathsf{EL}}, y\ \hat{\in}\ r, \exists r' \in o'.\ r \equiv_{\mathsf{Set}(T')} r'}{\Theta \vdash \Delta^{\mathsf{EL}}, w \subseteq_{T'} r, \exists r' \in o'.\ r \equiv_{\mathsf{Set}(T')} r'} \quad \forall\dfrac{\Theta, x \in r \vdash \Delta^{\mathsf{EL}}, x\ \hat{\in}\ w, \exists r' \in o'.\ r \equiv_{\mathsf{Set}(T')} r'}{\Theta \vdash \Delta^{\mathsf{EL}}, r \subseteq_{T'} w, \exists r' \in o'.\ r \equiv_{\mathsf{Set}(T')} r'}}{\exists\dfrac{\Theta \vdash \Delta^{\mathsf{EL}}, r \equiv_{\mathsf{Set}(T')} w, \exists r' \in o'.\ r \equiv_{\mathsf{Set}(T')} r'}{\Theta \vdash \Delta^{\mathsf{EL}}, \exists r' \in o'.\ r \equiv_{\mathsf{Set}(T')} r'}}$$

Applying the induction hypothesis to the leaves of that proof, we obtain two proofs of

$$\Theta, y \in w \vdash \Delta^{\mathsf{EL}}, y\ \hat{\in}\ r, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'$$

and

$$\Theta, x \in r \vdash \Delta^{\mathsf{EL}}, x\ \hat{\in}\ w, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'$$

so we can conclude using the admissibility of weakening and Corollary I.8 twice and replaying the $\wedge/\forall/\exists$ steps in the appropriate order (a branch of the proof derivation is elided below for lack of horizontal space):

$$\exists\dfrac{\forall\dfrac{\text{Lemma I.1}\dfrac{\vee\dfrac{\text{Corollary I.8}\dfrac{\Theta, y \in w \vdash \Delta^{\mathsf{EL}}, y\ \hat{\in}\ r, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'}{\Theta \vdash \Delta^{\mathsf{EL}}, \neg(y\ \hat{\in}\ w), y\ \hat{\in}\ r, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'}}{\Theta \vdash \Delta^{\mathsf{EL}}, y\ \hat{\in}\ w \to y\ \hat{\in}\ r, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'} \qquad \vdots}{\Theta \vdash \Delta^{\mathsf{EL}}, y\ \hat{\in}\ r \leftrightarrow y\ \hat{\in}\ w, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'}}{\Theta, y \in a \vdash \Delta^{\mathsf{EL}}, y\ \hat{\in}\ r \leftrightarrow y\ \hat{\in}\ w, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'}}{\Theta \vdash \Delta^{\mathsf{EL}}, \forall z \in a.\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ w, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'}}{\Theta \vdash \Delta^{\mathsf{EL}}, \exists r' \in o'.\forall z \in a,\ z\ \hat{\in}\ r \leftrightarrow z\ \hat{\in}\ r'}$$

□

Appendix J. Proof of the main theorem for non-set types

Recall again our main result:

Theorem 5.2 (Effective implicit to explicit for nested data). *Given a witness to the implicit definition of o in terms of $\vec{i}$ relative to Δ_0 $\varphi(\vec{i}, \vec{a}, o)$, one can compute NRC expression E such that for any $\vec{i}$, $\vec{a}$ and o, if $\varphi(\vec{i}, \vec{a}, o)$ then $E(\vec{i}) = o$. Furthermore, if the witness is EL-normalized, this can be done in polynomial time.*

In the body of the paper, we gave a proof for the case where the type of the defined object is $\mathsf{Set}(T)$ for any T. We now discuss the remaining cases: the base case and the inductive case for product types.

So assume we are given an implicit definition $\varphi(\vec{i}, \vec{a}, o)$ and a EL-normalized witness, and proceed by induction over the type o:

- If o has type Unit, then, since there is only one inhabitant in type Unit, then we can take our explicit definition to be the corresponding NRC expression $\langle\rangle$.
- If o has type $\mathfrak{U}$, then using interpolation on the entailment $\varphi(\vec{i}, \vec{a}, o) \to (\varphi(\vec{i}, \vec{a}', o') \to o =_{\mathfrak{U}} o')$, we obtain $\theta(\vec{i}, o)$ with $\varphi(\vec{i}, \vec{a}, o) \to \theta(\vec{i}, o)$ and $\theta(\vec{i}, o) \wedge \varphi(\vec{i}, \vec{a}, o') \to o =_{\mathfrak{U}} o'$. But then we know that φ implies o is a subobject of $\vec{i}$: otherwise we could find a model that contradicts the entailment. There is an NRC definition $A(\vec{i})$ that collects all of the $\mathfrak{U}$-elements lying beneath $\vec{i}$. We can then take $E(\vec{i}) = \text{GET}_T(\{x \in A(\vec{i}) \mid \theta(\vec{i}, x)\})$ as our NRC definition of o. The correctness of E follows from the properties of θ above.
- If o has type $T_1 \times T_2$, recalling the definition of $\equiv_{T_1 \times T_2}$, we have a derivation of

$$\varphi(\vec{i}, \vec{a}, o) \wedge \varphi(\vec{i}, \vec{a}, o') \vdash \pi_1(o) \equiv_{T_1} \pi_1(o') \wedge \pi_2(o) \equiv_{T_2} \pi_2(o')$$

By Lemma I.2, we have proofs of

$$\varphi(\vec{i}, \vec{a}, o) \wedge \varphi(\vec{i}, \vec{a}, o') \vdash \pi_i(o) \equiv_{T_i} \pi_i(o')$$

for $i \in \{1, 2\}$. Take o_1 and o_2 to be fresh variables of types T_1 and T_2. Take $\tilde{\varphi}(\vec{i}, \vec{a}, o_1, o_2)$ to be $\varphi(\vec{i}, \vec{a}, \langle o_1, o_2 \rangle)$. By substitutivity (the admissible rule given by Lemma I.4) and applying the $\times_\beta$ rule, we have EL-normalized proofs of

$$\varphi(\vec{i}, \vec{a}, \langle o_1, o_2 \rangle) \wedge \varphi(\vec{i}, \vec{a}, \langle o_1', o_2' \rangle) \vdash o_i \equiv_{T_i} o_i'$$

We can apply our inductive hypothesis to obtain a definition $E_i^{\mathsf{IH}}(\vec{i})$ for both $i \in \{1, 2\}$. We can then take our explicit definition to be $\langle E_1^{\mathsf{IH}}(\vec{i}), E_2^{\mathsf{IH}}(\vec{i}) \rangle$.

Appendix K. Variant of Parameter Collection Theorem, Theorem 5.9, for parameterized definability in first-order logic

Our paper has focused on the setting of nested relations, phrasing our results in terms of the language NRC. We indicated in the conclusion of the paper that there is a variant of the NRC parameter collection theorem, Theorem 5.9, for the broader context of first-order logic. In fact, this first-order version of the result provided the intuition for the theorem. In this section we present this variant.

We consider first-order logic with equality and without function symbols, which also excludes nullary function symbols, that is, individual constants, whose role is just taken by

free individual variables. Specifically, we consider first-order formulas with the following syntax

$$\varphi, \psi \ ::= \ P(\vec{x}) \mid \neg P(\vec{x}) \mid x = y \mid x \neq y \mid \top \mid \bot \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \forall x\, \varphi \mid \exists x\, \varphi.$$

Amongst the formulas given in the grammar above, those of the shape $P(\vec{x}), \neg P(\vec{x}), x = y$ and $x \neq y$ are called *literals*. Literals come with a polarity: those of the shape $P(\vec{x})$ and $x = y$ are positive while the other, of the shape $\neg P(\vec{x})$ or $x \neq y$ are negative.

On top of this, we give some "syntactic sugar". We define $\neg \varphi$ by induction over φ, dualizing every connective, including the quantifiers, and removing doubled negation over literals. We define implication $\varphi \to \psi$ as an abbreviation of $\neg \varphi \vee \psi$ and bi-implication $\varphi \leftrightarrow \psi$ as an abbreviation of $(\varphi \to \psi) \wedge (\psi \to \varphi)$. The set of free variables occurring in a formula φ is denoted by $FV(\varphi)$ and the set of predicates occurring in φ by $PR(\varphi)$.

Figure 5 shows our proof system for first-order logic. It is identical to a system from the prior literature.[4] Like the EL-normalized proof system we used in the body of the paper for Δ_0 formulas, it is a 1-sided calculus. The formulas other than Γ in the premise are the *active formulas* of the rule, while the *principal formulas* are the other formulas in its conclusion. The complementary principal formulas in Ax have to be literals. The replacement of symbols induced by equality with Repl is only performed on negative literals.

$$\mathrm{Ax}\ \frac{}{\vdash \Gamma, \varphi, \neg\varphi} \quad \varphi \ \textit{a positive literal} \qquad \top\ \frac{}{\vdash \Gamma, \top} \qquad \wedge\ \frac{\vdash \Gamma, \varphi_1 \qquad \vdash \Gamma, \varphi_2}{\vdash \Gamma, \varphi_1 \wedge \varphi_2}$$

$$\vee\ \frac{\vdash \Gamma, \varphi_1, \varphi_2}{\vdash \Gamma, \varphi_1 \vee \varphi_2}$$

$$\forall\ \frac{\vdash \Gamma, \varphi[y/x]}{\vdash \Gamma, \forall x\, \varphi} \quad y \notin FV(\Gamma, \forall x\, \varphi) \qquad \exists\ \frac{\vdash \Gamma, \varphi[t/x], \exists x\, \varphi}{\vdash \Gamma, \exists x\, \varphi}$$

$$\mathrm{Ref}\ \frac{\vdash t \neq t, \Gamma}{\vdash \Gamma} \qquad \mathrm{Repl}\ \frac{\vdash t \neq u, \varphi[u/x], \varphi[t/x], \Gamma}{\vdash t \neq u, \varphi[t/x], \Gamma} \quad \varphi \ \textit{a negative literal}$$

Figure 5: One-sided sequent calculus for first-order logic with equality.

As in the body of the paper, a proof tree or derivation is a tree whose nodes are labelled with sequents, such that the labels of the children of a given node and that of the node itself are the premises and conclusion, resp., of an instance of a rule from Figure 5. The *conclusion* of a proof tree is the sequent that labels its root. The proof system is closed under cut, weakening and contraction. Closure under contraction in particular makes it suited as basis for "root-first" proof search. Read in this "bottom-up" way, the ∃ rule states that a disjunction with an existentially quantified formula can be proven if the extension of the disjunction by a copy of the formula where the formerly quantified variable x is instantiated

---

[4]G3c+Ref+Repl [NvP01, TS00], in the one-sided form of GS3, discussed in Chapter 3 of [TS00], which reduces the number of rules.

with an arbitrary variable t can be proven. The existentially quantified formula is retained in the premise and may be used to add further instances by applying $\exists$ again in the course of the proof.

Soundness of the rules is straightforward. For example the $\exists$ rule could be read as stating that if we deduce a disjunction in which one disjunct is a formula φ with t in it, then we can deduce the same disjunction but with some occurrences of t replaced in that disjunct with an existentially quantified variable. Completeness of the proof system can also be proven by a standard Henkin-style construction: indeed, since this is really ordinary first-order logic, there are proofs in the literature for systems that are very similar to this one [TS00, NvP98].

Like our higher-level system in the body of the paper (Figure 2) the first-order system in Figure 5 does not impose any special constraints on the shape of the proof. We adapt the extra conditions on contexts of the EL-normalized system for Δ_0 formulas (Figure 3) to our first-order system. We characterize a proof in the system of Figure 5 as *FO-normalised* if no application of Ax, $\top$, $\exists$, REF, REPL contains in its conclusion a formula whose top-level connective is $\vee$, $\wedge$ or $\forall$.

The FO-normalised property may be either incorporated directly into a "root-first" proof procedure by constraining rule applications or it may be ensured by converting an arbitrary given proof tree to a FO-normalised proof tree with the same ultimate consequence. This conversion is achieved by a straightforward adaption of the method for the Δ_0 calculus from Appendix G, Figure 4. Only steps (1) and (2) of the method are relevant here, since step (3) is specifically for the pair terms of Δ_0 formulas. In step (2), instances of $\vee$, $\wedge$ or $\forall$ are permuted down over $\exists$, REF and REPL according to the generic permutation schemas of Lemma G.6. The left sides of $\vdash$ in these schemas can be ignored since they represent the $\in$-contexts for Δ_0 formulas. Finally, derivations of sequents containing $\top$ are replaced by applications of rule $\top$ and derivations of sequents containing complementary literals $\varphi, \neg\varphi$ by applications of rule Ax. The conversion, however, may increase the proof size exponentially as discussed in Section G.4.

We now discuss our generalization of the NRC Parameter Collection Theorem from the body of the paper to this first-order setting. The concept of explicit definition can be generalized to *definition up to parameters and disjunction*: A family of formulas $\chi_i(\vec{z}, \vec{y}, \vec{r})$, $1 \leq i \leq n$, provides an *explicit definition up to parameters and disjunction* of a formula $\lambda(\vec{z}, \vec{l})$ relative to a formula φ if

$$\varphi \models \bigvee_{i=1}^{n} \exists\vec{y}\forall\vec{z}\,(\lambda(\vec{z}, \vec{l}) \leftrightarrow \chi_i(\vec{z}, \vec{y}, \vec{r})). \tag{$\star$}$$

The entailment ($\star$) is considered with restrictions on the predicates and variables permitted to occur in the χ_i. In the simplest case, $\lambda(\vec{z}, \vec{l})$ is a positive literal $p(\vec{z})$ with a predicate that is permitted in φ but not in the χ_i. The predicate p is then said to be *explicitly definable up to parameters and disjunction* with respect to φ [CK92].

The disjunction over a finite family of formulas χ_i can be consolidated into a single quantified biconditional as long as the domain has size at least 2 in every model of φ. Notice that if φ has only finite models, then by the compactness theorem of first-order logic, the size of models must be bounded. In such cases every formula λ is definable with sufficiently many parameters.

We can now state our analog of the Parameter Collection Theorem, Theorem 5.9.

Theorem K.1. *Let φ, ψ, $\lambda(\vec{z}, \vec{l})$, and $\rho(\vec{z}, \vec{y}, \vec{r})$ be first-order formulas such that*

$$\varphi \wedge \psi \models \exists \vec{y} \forall \vec{z} (\lambda(\vec{z}, \vec{l}) \leftrightarrow \rho(\vec{z}, \vec{y}, \vec{r})).$$

Then there exist first-order formulas $\chi_i(\vec{z}, \vec{v}_i, \vec{c}_i)$, $1 \leq i \leq n$, such that

(1) $\varphi \wedge \psi \models \bigvee_{i=1}^{n} \exists \vec{v}_i \forall \vec{z} (\lambda(\vec{z}, \vec{l}) \leftrightarrow \chi_i(\vec{z}, \vec{v}_i, \vec{c}_i))$,

(2) $\vec{c}_i \subseteq (FV(\varphi) \cup \vec{l}) \cap (FV(\psi) \cup \vec{r})$,

(3) $PR(\chi_i) \subseteq (PR(\varphi) \cup PR(\lambda)) \cap (PR(\psi) \cup PR(\rho))$.

Moreover, given a FO-normalised proof of the precondition with the system of Figure 5, a family of formulas χ_i, $1 \leq i \leq n$, with the claimed properties can be computed in polynomial time in the size of the proof tree.

In the theorem statement, the free variables of λ are $\vec{z}$ and $\vec{l}$, and the free variables of ρ are $\vec{z}$, $\vec{y}$, and $\vec{r}$. The precondition supposes an explicit definition ρ of λ up to parameters with respect to a conjunction $\varphi \wedge \psi$. The conclusion then claims that one can effectively compute another definition of λ with respect to $\varphi \wedge \psi$ that is up to parameters and disjunction and has a constrained signature: free variables and predicates must occur in at least one of the "left side" formulas φ and λ and also in at least one of the "right side" formulas ψ and ρ. In other words, the theorem states that if $\mathcal{SIG}_L$ and $\mathcal{SIG}_R$ are "left" and "right" signatures such that φ and λ are over $\mathcal{SIG}_L$, ψ and ρ are over $\mathcal{SIG}_R$, and ρ provides an explicit definition of λ up to parameters with respect to $\varphi \wedge \psi$, then one can effectively compute another definition of λ with respect to $\varphi \wedge \psi$ that is up to parameters and disjunction and is just over the intersection of the signatures $\mathcal{SIG}_L$ and $\mathcal{SIG}_R$.

We now prove Theorem K.1 by induction on the depth of the proof tree, generalizing the constructive proof method for Craig interpolation often called Maehara's method [Tak87, TS00, Smu68b]. To simplify the presentation we assume that the tuples $\vec{z}$ and $\vec{y}$ in the theorem statement each consist of a single variable z and y, respectively. The generalization of our argument to tuples of variables is straightforward.

To specify conveniently the construction steps of the family of formulas χ_i we introduce the following concept: A *pre-defining equivalence up to parameters and disjunction* (briefly PDEPD) for a formula $\lambda(z, \vec{l})$ is a formula δ built up from formulas of the form $\forall z\,(\lambda(z, \vec{l}) \leftrightarrow \chi(z, \vec{p}))$ (where the left side is always the same formula $\lambda(z, \vec{l})$ but the right sides $\chi(z, \vec{p})$ may differ) and a finite number of applications of disjunction and existential quantification upon variables from the vectors $\vec{p}$ of the right sides. The empty disjunction $\bot$ is allowed as a special case of a PDEPD. By rewriting with the equivalence $\exists v\,(\delta_1 \vee \delta_2) \equiv \exists v\,\delta_1 \vee \exists v\,\delta_2$, any PDEPD for λ can be efficiently transformed into the form $\bigvee_{i=1}^{n} \exists \vec{v}_i \forall z\,(\lambda(z, \vec{l}) \leftrightarrow \chi_i(z, \vec{v}_i, \vec{r}_i))$ for some natural number $n \geq 0$. That is, although a PDEPD has in general not the syntactic form of the disjunction of quantified biconditionals in the theorem statement (thus "pre-"), it corresponds to such a disjunction. The more generous syntax will be convenient in the induction. The sets of additional variables $\vec{l}$ and $\vec{p}$ in the biconditionals $\lambda(z, \vec{l}) \leftrightarrow \chi(z, \vec{p})$ can overlap, but the overlap will be top-level variables that never get quantified. Although we have defined the notion of PDEPD for a general λ, in the proof we just consider PDEPDs for the formula λ from the theorem statement.

For PDEPDs δ we provide analogs to FV and PR that only yield free variables and predicates of δ that occur in a right side of its biconditionals, which helps to express the restrictions by definability properties that constrains the signature of exactly those right sides. Recall that we refer of these right sides as subformulas $\chi(z, \vec{p})$. For PDEPD δ, define $PR^{\mathsf{RHS}}(\delta)$ as the set of the predicate symbols that occur in a subformula $\chi(z, \vec{p})$ of δ and

define $FV^{\mathsf{RHS}}(\delta)$ as the set of all variables that occur in a subformula $\chi(z,\vec{p})$ of δ and are free in δ. In other words, $FV^{\mathsf{RHS}}(\delta)$ is the set of all variables p in the vectors $\vec{p}$ of the subformulas $\chi(z,\vec{p})$ that have an occurrence in δ which is not in the scope of a quantifier $\exists p$. If, for example $\delta = \bigvee_{i=1}^{n} \exists \vec{v_i} \forall z\, (\lambda(z,\vec{l}) \leftrightarrow \chi_i(z,\vec{v_i},\vec{r_i}))$, then $FV^{\mathsf{RHS}}(\delta)$ is the set of all variables in the vectors r_i, for $1 \leq i \leq n$.

To build up PDEPDs we provide an operation that only affects the right sides of the biconditionals, a restricted form of existential quantification. It is for use in interpolant construction to convert a free variable in right sides that became illegal into an existentially quantified parameter. For variables p, y define $\delta[y/p]^{\mathsf{RHS}}$ as δ after substituting all occurrences of p that are within a right side $\chi(z,\vec{p})$ and are free in δ with y. Define $\exists^{\mathsf{RHS}} p\, \delta$ as shorthand for $\exists y\, \delta[y/p]^{\mathsf{RHS}}$, where y is a fresh variable. Clearly $\delta \models \exists^{\mathsf{RHS}} p\, \delta$ and p has no free occurrences in $\exists^{\mathsf{RHS}} p\, \delta$ that are in any of the right side formulas $\chi(z,\vec{p})$, i.e., $p \notin FV^{\mathsf{RHS}}(\exists^{\mathsf{RHS}} p\, \delta)$. Occurrences of p in the left sides $\lambda(z,\vec{l})$, if p is a member of $\vec{l}$, are untouched in $\exists^{\mathsf{RHS}} p\, \delta$. If p is not in $\vec{l}$, then $\exists^{\mathsf{RHS}} p\, \delta$ reduces to ordinary existential quantification $\exists y\, \delta[y/p]$.

We introduce the following symbolic shorthand for the parametric definition on the right side in the theorem's precondition.

$$\mathcal{G} \quad := \quad \exists y \forall z\, (\lambda(z,\vec{l}) \leftrightarrow \rho(z,y,\vec{r})).$$

Note that our proof rules are such that if we have a proof (FO-normalised or not) that our original top-level "global" parametric definition is implied by some formula, then every one-sided sequent in the proof must include that parametric definition in it. This is because the rules that eliminate a formula when read "bottom-up" cannot apply to that parametric definition, whose outermost logic operator is the existential quantifier. Thus, in our inductive argument, we can assume that $\mathcal{G}$ is always present.

We write

$$\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle,$$

where $\vdash \Gamma_L, \Gamma_R, \mathcal{G}$ is a sequent, partitioned into three components, multisets Γ_L and Γ_R of formulas and the formula $\mathcal{G}$ from the theorem's hypothesis, θ is a formula and $\mathcal{D}$ is a PDEPD, to express that the following properties hold:

I1. $\models \Gamma_R \vee \theta$.
I2. $\models \neg\theta \vee \Gamma_L \vee \mathcal{D}$.
I3. $PR(\theta) \subseteq (PR(\Gamma_L) \cup PR(\lambda)) \cap (PR(\Gamma_R) \cup PR(\rho))$.
I4. $FV(\theta) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.
I5. $\mathcal{D}$ is a PDEPD for $\lambda(z,\vec{l})$.
I6. $PR^{\mathsf{RHS}}(\mathcal{D}) \subseteq (PR(\Gamma_L) \cup PR(\lambda)) \cap (PR(\Gamma_R) \cup PR(\rho))$.
I7. $FV^{\mathsf{RHS}}(\mathcal{D}) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.

For a given proof with conclusion $\vdash \neg\varphi, \neg\psi, \mathcal{G}$, corresponding to the hypothesis $\varphi \wedge \psi \models \mathcal{G}$ of the theorem, we show the construction of a formula θ and PDEPD $\mathcal{D}$ such that

$$\vdash \neg\varphi;\ \neg\psi;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle.$$

From properties I1.–I2. and I5.– I7. it is then straightforward to read off that the formula $\bigvee_{i=1}^{n} \exists \vec{v_i} \forall \vec{z} (\lambda(\vec{z},\vec{l}) \leftrightarrow \chi_i(\vec{z},\vec{v_i},\vec{r_i}))$ obtained from $\mathcal{D}$ by propagating existential quantifiers inwards is as claimed in the theorem's conclusion.

Formula θ plays an auxiliary role in the induction. For the overall conclusion of the proof it a side result that is like a Craig interpolant of ψ and $\varphi \to D$, but slightly weaker

syntactically constrained by taking λ and ρ into account: $FV(\theta) \subseteq ((FV(\varphi) \cup \vec{l}) \cap (FV(\psi) \cup \vec{r}))$ and $PR(\theta) \subseteq (PR(\varphi) \cup PR(\lambda)) \cap (PR(\psi) \cup PR(\rho))$.

As basis of the induction, we have to show constructions of θ and $\mathcal{D}$ such that $\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G} : \langle \theta, \mathcal{D} \rangle$ holds for Ax and $\top$, considering each possibility in which the principal formula(s) can be in Γ_L or Γ_R. For the induction step, there are a number of subcases, according to which rule is last applied and which of the partitions Γ_L, Γ_R or $\mathcal{G}$ contain the principal formula(s). We first discuss the most interesting case, the induction step where $\mathcal{G}$ is the principal formula. This case is similar to the most interesting case in the NRC Parameter Collection Theorem, covered in the body of the paper.

Case where the principal formula is $\mathcal{G}$. We now give more detail on the most complex case. If the principal formula of a conclusion $\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G}$ is $\mathcal{G}$, then the rule that is applied rule must be $\exists$. From the FO-normalised property of the proof it follows that the derivation tree ending in $\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G}$ must have the following shape, for some $u \notin FV(\Gamma_L, \Gamma_R, \mathcal{G})$ and $w \neq u$. Note that u could be either a top-level variable from λ, i.e., a member of $\vec{l}$, or one introduced during the proof.

$$
\begin{array}{c}
\wedge \dfrac{\vee \dfrac{\vdash \Gamma_L, \neg\lambda(u,\vec{l}), \Gamma_R, \rho(u,w,\vec{r}), \mathcal{G}}{\vdash \Gamma_L, \Gamma_R, \lambda(u,\vec{l}) \to \rho(u,w,\vec{r}), \mathcal{G}} \qquad \vee \dfrac{\vdash \Gamma_L, \lambda(u,\vec{l}), \Gamma_R, \neg\rho(u,w,\vec{r}), \mathcal{G}}{\vdash \Gamma_L, \lambda(u,\vec{l}), \Gamma_R, \rho(u,w,\vec{r}) \to \lambda(u,\vec{l}), \mathcal{G}}}{\exists \dfrac{\forall \dfrac{\vdash \Gamma_L, \Gamma_R, \lambda(u,\vec{l}) \leftrightarrow \rho(u,w,\vec{r}), \mathcal{G}}{\vdash \Gamma_L, \Gamma_R, \forall z\,(\lambda(z,\vec{l}) \leftrightarrow \rho(z,w,\vec{r})), \mathcal{G}}}{\vdash \Gamma_L, \Gamma_R, \mathcal{G}}}
\end{array}
$$

The important point is that the two "leaves" of the above tree are both sequents where we can apply our induction hypothesis. Taking into account the partitioning of the sequents at the bottom conclusion and the top premises in this figure, we can express the induction step in the form of a "macro" rule that specifies the how we constructed the required θ and $\mathcal{D}$ for the conclusion, making use of the $\theta_1, \mathcal{D}_1$ and $\theta_2, \mathcal{D}_2$ that we get by applying the induction hypothesis to each of the two premises.

$$
\dfrac{\vdash \Gamma_L, \neg\lambda(u,\vec{l});\ \Gamma_R, \rho(u,w,\vec{r});\ \mathcal{G} : \langle \theta_1, \mathcal{D}_1 \rangle \quad \vdash \Gamma_L, \lambda(u,\vec{l});\ \Gamma_R, \neg\rho(u,w,\vec{r});\ \mathcal{G} : \langle \theta_2, \mathcal{D}_2 \rangle}{\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G} : \langle \theta, \mathcal{D} \rangle,}
$$

where u is as above. The values of θ and $\mathcal{D}$ – the new formula and definition that we are building – will depend on occurrences of w, and we give their construction in cases below:

(i) If $w \notin FV(\Gamma_L) \cup \vec{l}$ or $w \in FV(\Gamma_R) \cup \vec{r}$, then

$$
\begin{aligned}
\theta &:= \forall u\,(\theta_1 \vee \theta_2). \\
\mathcal{D} &:= \exists^{\mathsf{RHS}} u\,\mathcal{D}_1 \vee \exists^{\mathsf{RHS}} u\,\mathcal{D}_2 \vee \forall z\,(\lambda(z,\vec{l}) \leftrightarrow \theta_2[z/u]).
\end{aligned}
$$

(ii) Else it holds that $w \in FV(\Gamma_L) \cup \vec{l}$ and $w \notin FV(\Gamma_R) \cup \vec{r}$. Then

$$
\begin{aligned}
\theta &:= \forall w \forall u\,(\theta_1 \vee \theta_2). \\
\mathcal{D} &:= \exists^{\mathsf{RHS}} w\,(\exists^{\mathsf{RHS}} u\,\mathcal{D}_1 \vee \exists^{\mathsf{RHS}} u\,\mathcal{D}_2 \vee \forall z\,(\lambda(z,\vec{l}) \leftrightarrow \theta_2[z/u])).
\end{aligned}
$$

We now verify that $\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G} : \langle \theta', \mathcal{D}' \rangle$, that is, properties I1.–I7., hold. The proofs for the individual properties are presented in tabular form, with explanations annotated in the side column, where *IH* stands for *induction hypothesis*. We concentrate on the case (i) and indicate the modifications of the proofs for case (ii) in remarks, where we refer to the

values of θ and $\mathcal{D}$ for that case in terms of the values for the case (i) as $\forall w\,\theta$ and $\exists^{\mathsf{RHS}}w\,\mathcal{D}$. In the proofs of the semantic properties I1. and I2. we let sequents stand for the disjunction of their members.

Property I1.:

$$
\begin{array}{lll}
(1) & \models \Gamma_R \vee \rho(u,w,\vec{r}) \vee \theta_1. & \text{IH} \\
(2) & \models \Gamma_R \vee \neg\rho(u,w,\vec{r}) \vee \theta_2. & \text{IH} \\
(3) & \models \Gamma_R \vee \theta_1 \vee \theta_2. & \text{by (1), (2)} \\
(4) & \models \Gamma_R \vee \forall u\,(\theta_1 \vee \theta_1) & \text{by (3) since } u \notin FV(\Gamma_R) \\
(5) & \models \Gamma_R \vee \theta. & \text{by (4) and def. of } \theta
\end{array}
$$

For case (ii), it follows from the precondition $w \notin FV(\Gamma_R)$ and step (5) that $\models \Gamma_R \vee \forall w\,\theta$.

Property I2.:

$$
\begin{array}{lll}
(1) & \models \neg\theta_1 \vee \Gamma_L \vee \neg\lambda(u,\vec{l}) \vee \mathcal{D}_1. & \text{IH} \\
(2) & \models \neg\theta_2 \vee \Gamma_L \vee \lambda(u,\vec{l}) \vee \mathcal{D}_2. & \text{IH} \\
(3) & \models \neg(\theta_1 \vee \theta_2) \vee \Gamma_L \vee \neg\lambda(u,\vec{l}) \vee \mathcal{D}_1 \vee \theta_2. & \text{by (1)} \\
(4) & \models \neg(\theta_1 \vee \theta_2) \vee \Gamma_L \vee \neg\lambda(u,\vec{l}) \vee \mathcal{D}_1 \vee \mathcal{D}_2 \vee \theta_2. & \text{by (3)} \\
(5) & \models \neg\forall u\,(\theta_1 \vee \theta_2) \vee \Gamma_L \vee \neg\lambda(u,\vec{l}) \vee \mathcal{D}_1 \vee \mathcal{D}_2 \vee \theta_2. & \text{by (4)} \\
(6) & \models \neg\forall u\,(\theta_1 \vee \theta_2) \vee \neg\theta_2 \vee \Gamma_L \vee \lambda(u,\vec{l}) \vee \mathcal{D}_1 \vee \mathcal{D}_2. & \text{by (2)} \\
(7) & \models \neg\forall u\,(\theta_1 \vee \theta_2) \vee \Gamma_L \vee \mathcal{D}_1 \vee \mathcal{D}_2 \vee & \\
& \quad (\lambda(u,\vec{l}) \leftrightarrow \theta_2). & \text{by (6), (5)} \\
(8) & \models \neg\forall u\,(\theta_1 \vee \theta_2) \vee \Gamma_L \vee \exists^{\mathsf{RHS}}u\,\mathcal{D}_1 \vee \exists^{\mathsf{RHS}}u\,\mathcal{D}_2 \vee & \\
& \quad (\lambda(u,\vec{l}) \leftrightarrow \theta_2). & \text{by (7)} \\
(9) & \models \neg\forall u\,(\theta_1 \vee \theta_2) \vee \Gamma_L \vee \exists^{\mathsf{RHS}}u\,\mathcal{D}_1 \vee \exists^{\mathsf{RHS}}u\,\mathcal{D}_2 \vee & \\
& \quad \forall z\,(\lambda(z,\vec{l}) \leftrightarrow \theta_2[z/u]). & \text{by (8) since } u \notin FV(\Gamma_L, \lambda(z,\vec{l})) \\
(10) & \models \neg\theta \vee \Gamma_L \vee \mathcal{D}. & \text{by (9), defs. } \theta, \mathcal{D}
\end{array}
$$

That $u \notin FV(\lambda(z,\vec{l}))$ follows from the precondition $u \notin FV(\mathcal{G})$. It is used in step (9) to justify that the substitution $[z/u]$ has only to be applied to θ_2 and not to $\lambda(z,\vec{l})$ and, in addition, to justify that $u \notin FV(\exists^{\mathsf{RHS}}u\,\mathcal{D}_1)$ and $u \notin FV(\exists^{\mathsf{RHS}}u\,\mathcal{D}_2)$, which follow from $u \notin FV(\lambda(z,\vec{l}))$ and the induction hypotheses that property I5. applies to $\mathcal{D}_1$ and $\mathcal{D}_2$.

For case (ii), it follows from step (10) that $\models \neg\forall w\,\theta \vee \Gamma_L \vee \exists^{\mathsf{RHS}}w\,\mathcal{D}$.

Property I3.:

$$
\begin{array}{lll}
(1) & PR(\theta_1) \subseteq (PR(\Gamma_L, \neg\lambda(u,\vec{l})) \cup PR(\lambda)) \cap (PR(\Gamma_R, \rho(u,w,\vec{r})) \cup PR(\rho)). & \text{IH} \\
(2) & PR(\theta_2) \subseteq (PR(\Gamma_L, \lambda(u,\vec{l})) \cup PR(\lambda)) \cap (PR(\Gamma_R, \neg\rho(u,w,\vec{r})) \cup PR(\rho)). & \text{IH} \\
(3) & PR(\forall u\,(\theta_1 \vee \theta_1)) & \text{by (1), (2)} \\
(4) & PR(\theta) \subseteq (PR(\Gamma_L) \cup PR(\lambda)) \cap (PR(\Gamma_R) \cup PR(\rho)) & \text{by (3), def. } \theta
\end{array}
$$

For case (ii) the property follows since $PR(\theta) = PR(\forall w\,\theta)$.

Property I4.:

(1) $FV(\theta_1) \subseteq (FV(\Gamma_L, \neg\lambda(u, \vec{l})) \cup \vec{l}) \cap (FV(\Gamma_R, \rho(u, w, \vec{r})) \cup \vec{r})$. IH

(2) $FV(\theta_2) \subseteq (FV(\Gamma_L, \lambda(u, \vec{l})) \cup \vec{l}) \cap (FV(\Gamma_R, \neg\rho(u, w, \vec{r})) \cup \vec{r})$. IH

(3) $FV(\theta_1) \subseteq (FV(\Gamma_L) \cup \vec{l} \cup \{u\}) \cap (FV(\Gamma_R) \cup \vec{r} \cup \{u, w\})$.   by (1)

(4) $FV(\theta_1) \subseteq (FV(\Gamma_L) \cup \vec{l} \cup \{u\}) \cap (FV(\Gamma_R) \cup \vec{r} \cup \{u, w\})$.   by (2)

(5) $FV(\forall u\, (\theta_1 \vee \theta_2)) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.   by (3), (4),

and $w \notin FV(\Gamma_L) \cup \vec{l}$ or $w \in FV(\Gamma_R) \cup \vec{r}$

(6) $FV(\theta) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.   by (5), def. θ

For case (ii) step (5) has to be replaced by

$$FV(\forall w\forall u\, (\theta_1 \vee \theta_2)) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r}),$$

which follows just from (3) and (4). Instead of step (6) we then have $FV(\forall w\, \theta) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.

Property I5.: Immediate from the induction hypothesis and the definition of $\mathcal{D}$.

Property I6.:

(1) $PR^{\mathsf{RHS}}(\mathcal{D}_1) \subseteq (PR(\Gamma_L, \neg\lambda(u, \vec{l})) \cup PR(\lambda)) \cap (PR(\Gamma_R, \rho(u, w, \vec{r})) \cup PR(\rho))$. IH

(2) $PR^{\mathsf{RHS}}(\mathcal{D}_2) \subseteq (PR(\Gamma_L, \lambda(u, \vec{l})) \cup PR(\lambda)) \cap (PR(\Gamma_R, \neg\rho(u, w, \vec{r})) \cup PR(\rho))$. IH

(3) $PR(\theta_2) \subseteq (PR(\Gamma_L, \lambda(u, \vec{l})) \cup PR(\lambda)) \cap (PR(\Gamma_R, \neg\rho(u, w, \vec{r})) \cup PR(\rho))$.   IH

(4) $PR^{\mathsf{RHS}}(\exists^{\mathsf{RHS}}u\, \mathcal{D}_1 \vee \exists^{\mathsf{RHS}}u\, \mathcal{D}_2 \vee \forall z\, (\lambda(z, \vec{l}) \leftrightarrow \theta_2[z/u])) \subseteq$
$(PR(\Gamma_L) \cup PR(\lambda)) \cap (PR(\Gamma_R) \cup PR(\rho))$ by (1)–(3)

(5) $PR^{\mathsf{RHS}}(\mathcal{D}) \subseteq (PR(\Gamma_L) \cup PR(\lambda)) \cap (PR(\Gamma_R) \cup PR(\rho))$ by (4), def. $\mathcal{D}$

For case (ii) the property follows since $PR^{\mathsf{RHS}}(\mathcal{D}) = PR^{\mathsf{RHS}}(\exists^{\mathsf{RHS}}w\, \mathcal{D})$.

Property I7.:

(1) $FV^{\mathsf{RHS}}(\mathcal{D}_1) \subseteq (FV(\Gamma_L, \neg\lambda(u, \vec{l})) \cup \vec{l}) \cap (FV(\Gamma_R, \rho(u, w, \vec{r})) \cup \vec{r})$. IH

(2) $FV^{\mathsf{RHS}}(\mathcal{D}_2) \subseteq (FV(\Gamma_L, \lambda(u, \vec{l})) \cup \vec{l}) \cap (FV(\Gamma_R, \neg\rho(u, w, \vec{r})) \cup \vec{r})$. IH

(3) $FV(\theta_2) \subseteq (FV(\Gamma_L, \lambda(u, \vec{l})) \cup \vec{l}) \cap (FV(\Gamma_R, \neg\rho(u, w, \vec{r})) \cup \vec{r})$.   IH

(4) $FV^{\mathsf{RHS}}(\mathcal{D}_1) \subseteq (FV(\Gamma_L) \cup \vec{l} \cup \{u\}) \cap (FV(\Gamma_R) \cup \vec{r} \cup \{u, w\})$.   by (1)

(5) $FV^{\mathsf{RHS}}(\mathcal{D}_2) \subseteq (FV(\Gamma_L) \cup \vec{l} \cup \{u\}) \cap (FV(\Gamma_R) \cup \vec{r} \cup \{u, w\})$.   by (2)

(6) $FV(\theta_2) \subseteq (FV(\Gamma_L) \cup \vec{l} \cup \{u\}) \cap (FV(\Gamma_R) \cup \vec{r} \cup \{u, w\})$.   by (3)

(7) $FV^{\mathsf{RHS}}(\exists^{\mathsf{RHS}}u\, \mathcal{D}_1 \vee \exists^{\mathsf{RHS}}u\, \mathcal{D}_2 \vee \forall z\, (\lambda(z, \vec{l}) \leftrightarrow \theta_2[z/u])) \subseteq$
$(FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.   by (4), (5), (6),

and $w \notin FV(\Gamma_L) \cup \vec{l}$ or $w \in FV(\Gamma_R) \cup \vec{r}$

(8) $FV^{\mathsf{RHS}}(\mathcal{D}) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.   by (7), def. $\mathcal{D}$

For case (ii) step (7) has to be replaced by

$$FV^{\mathsf{RHS}}(\exists^{\mathsf{RHS}}w(\exists^{\mathsf{RHS}}u\, \mathcal{D}_1 \vee \exists^{\mathsf{RHS}}u\, \mathcal{D}_2 \vee \forall z\, (\lambda(z, \vec{l}) \leftrightarrow \theta_2[z/u]))) \subseteq$$
$$(FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r}),$$

which follows just from (4), (5), (6). Instead of step (8) we then have $FV^{\mathsf{RHS}}(\exists^{\mathsf{RHS}}w\, \mathcal{D}) \subseteq (FV(\Gamma_L) \cup \vec{l}) \cap (FV(\Gamma_R) \cup \vec{r})$.

This completes the verification of correctness, and thus ends our discussion of this case. We now turn to the base of the induction along with the other inductive cases.

Cases where the principal formulas are in the Γ_L or Γ_R partition. The inductive cases where the principal formulas are in the Γ_L or Γ_R partition can be conveniently specified in the form of rules that lead from induction hypotheses of the form $\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta, \mathcal{D}\rangle$ as premises to an induction conclusion of the same form. Base cases can be taken as rules without premises. The axioms and rules shown below correspond to those of the calculus, but replicated for each possible way in which the partitions Γ_L, Γ_R or $\mathcal{G}$ of the conclusion can contain the principal formula(s). To verify that properties I1.–I7. are preserved by each of the constructions is straightforward, and therefore we only point out a few subtleties.

(1) Ax $\dfrac{}{\vdash \Gamma_L, \varphi, \neg\varphi;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\top, \bot\rangle}$ *φ a positive literal*

(2) Ax $\dfrac{}{\vdash \Gamma_L;\ \Gamma_R, \varphi, \neg\varphi;\ \mathcal{G}\ :\ \langle\bot, \bot\rangle}$ *φ a positive literal*

(3) Ax $\dfrac{}{\vdash \Gamma_L, \varphi;\ \Gamma_R, \neg\varphi;\ \mathcal{G}\ :\ \langle\varphi, \bot\rangle}$ *φ a literal*

(4) $\top$ $\dfrac{}{\vdash \Gamma_L, \top;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\top, \bot\rangle}$

(5) $\top$ $\dfrac{}{\vdash \Gamma_L;\ \Gamma_R, \top;\ \mathcal{G}\ :\ \langle\bot, \bot\rangle}$

(6) $\vee$ $\dfrac{\vdash \Gamma_L, \varphi_1, \varphi_2;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta, \mathcal{D}\rangle}{\vdash \Gamma_L, \varphi_1 \vee \varphi_2;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta, \mathcal{D}\rangle}$

(7) $\vee$ $\dfrac{\vdash \Gamma_L;\ \Gamma_R, \varphi_1, \varphi_2;\ \mathcal{G}\ :\ \langle\theta, \mathcal{D}\rangle}{\vdash \Gamma_L;\ \Gamma_R, \varphi_1 \vee \varphi_2;\ \mathcal{G}\ :\ \langle\theta, \mathcal{D}\rangle}$

(8) $\wedge$ $\dfrac{\vdash \Gamma_L, \varphi_1;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta_1, \mathcal{D}_1\rangle \qquad \vdash \Gamma_L, \varphi_2;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta_2, \mathcal{D}_2\rangle}{\vdash \Gamma_L, \varphi_1 \wedge \varphi_2;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta_1 \wedge \theta_2, \mathcal{D}_1 \vee \mathcal{D}_2\rangle}$

(9) $\wedge$ $\dfrac{\vdash \Gamma_L;\ \Gamma_R, \varphi_1;\ \mathcal{G}\ :\ \langle\theta_1, \mathcal{D}_1\rangle \qquad \vdash \Gamma_L;\ \Gamma_R, \varphi_2;\ \mathcal{G}\ :\ \langle\theta_2, \mathcal{D}_2\rangle}{\vdash \Gamma_L;\ \Gamma_R, \varphi_1 \wedge \varphi_2;\ \mathcal{G}\ :\ \langle\theta_1 \vee \theta_2, \mathcal{D}_1 \vee \mathcal{D}_2\rangle}$

(10) $\exists$ $\dfrac{\vdash \Gamma_L, \varphi[t/x], \exists x\, \varphi;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta, \mathcal{D}\rangle}{\vdash \Gamma_L, \exists x\, \varphi;\ \Gamma_R;\ \mathcal{G}\ :\ \langle\theta', \mathcal{D}'\rangle,}$
where the values of θ' and $\mathcal{D}'$ depend on occurrences of t:
- If $t \in FV(\Gamma_L, \exists x\, \varphi) \cup \vec{l}$, then $\theta' := \theta$ and $\mathcal{D}' := \mathcal{D}$.
- Else it holds that $t \notin FV(\Gamma_L, \exists x\, \varphi) \cup \vec{l}$. Then $\theta' := \exists t\, \theta$ and $\mathcal{D}' := \exists^{\mathsf{RHS}} t\, \mathcal{D}$.

(11) $\exists$ $\dfrac{\vdash \Gamma_L;\ \Gamma_R, \varphi[t/x], \exists x\, \varphi;\ \mathcal{G}\ :\ \langle\theta, \mathcal{D}\rangle}{\vdash \Gamma_L;\ \Gamma_R, \exists x\, \varphi;\ \mathcal{G}\ :\ \langle\theta', \mathcal{D}'\rangle,}$
where the values of θ' and $\mathcal{D}'$ depend on occurrences of t:
- If $t \in FV(\Gamma_R, \exists x\, \varphi) \cup \vec{r}$, then $\theta' := \theta$ and $\mathcal{D}' := \mathcal{D}$.
- Else it holds that $t \notin FV(\Gamma_R, \exists x\, \varphi) \cup \vec{r}$. Then $\theta' := \forall t\, \theta$ and $\mathcal{D}' := \exists^{\mathsf{RHS}} t\, \mathcal{D}$.

(12) $\quad \forall \dfrac{\vdash \Gamma_L, \varphi[y/x];\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash \Gamma_L, \forall x\, \varphi;\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle} \qquad y \notin FV(\Gamma_L, \forall x\, \varphi, \Gamma_R, \mathcal{G})$

(13) $\quad \forall \dfrac{\vdash \Gamma_L;\ \Gamma_R, \varphi[y/x];\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash \Gamma_L;\ \Gamma_R, \forall x\, \varphi;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle} \qquad y \notin FV(\Gamma_L, \Gamma_R, \forall x\, \varphi, \mathcal{G})$

(14) $\quad \text{Ref} \dfrac{\vdash t \neq t, \Gamma_L;\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle} \qquad t \in FV(\Gamma_L) \cup \vec{l}$

(15) $\quad \text{Ref} \dfrac{\vdash \Gamma_L;\ t \neq t, \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash \Gamma_L;\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle} \qquad t \notin FV(\Gamma_L) \cup \vec{l}$

(16) $\quad \text{Repl} \dfrac{\vdash t \neq u, \varphi[u/x], \varphi[t/x], \Gamma_L;\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash t \neq u, \varphi[t/x], \Gamma_L;\ \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle} \qquad \varphi\ \text{a negative literal}$

(17) $\quad \text{Repl} \dfrac{\vdash t \neq u, \Gamma_L;\ \varphi[u/x], \varphi[t/x], \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash t \neq u, \Gamma_L;\ \varphi[t/x], \Gamma_R;\ \mathcal{G}\ :\ \langle \theta', \mathcal{D}' \rangle,} \qquad \varphi\ \text{a negative literal}$

where the values of θ' and $\mathcal{D}'$ depend on occurrences of t and u:
- If $t \notin FV(\varphi[t/x], \Gamma_R) \cup \vec{r}$, then $\theta' := \theta$ and $\mathcal{D}' := \mathcal{D}$. In this subcase the precondition $t \notin FV(\varphi[t/x])$ implies that $x \notin FV(\varphi)$ and thus $\varphi[u/x] = \varphi[t/x]$.
- If $t, u \in FV(\varphi[t/x], \Gamma_R) \cup \vec{r}$, then $\theta' := \theta \vee t \neq u$ and $\mathcal{D}' := \mathcal{D}$.
- Else it holds that $t \in FV(\varphi[t/x], \Gamma_R) \cup \vec{r}$ and $u \notin FV(\varphi[t/x], \Gamma_R) \cup \vec{r}$. Then $\theta' := \theta[t/u]$ and $\mathcal{D}' := \mathcal{D}[t/u]^{\mathsf{RHS}}$. For this subcase, to derive property I1. it is used that the precondition $u \notin \varphi[t/x]$ implies that $\varphi[u/x][t/u] = \varphi[t/x]$.

(18) $\quad \text{Repl} \dfrac{\vdash \Gamma_L;\ t \neq u, \varphi[u/x], \varphi[t/x], \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash \Gamma_L;\ t \neq u, \varphi[t/x], \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle} \qquad \varphi\ \text{a negative literal}$

(19) $\quad \text{Repl} \dfrac{\vdash \varphi[u/x], \varphi[t/x], \Gamma_L;\ t \neq u; \Gamma_R;\ \mathcal{G}\ :\ \langle \theta, \mathcal{D} \rangle}{\vdash \varphi[t/x], \Gamma_L;\ t \neq u, \Gamma_R;\ \mathcal{G}\ :\ \langle \theta', \mathcal{D}' \rangle,} \qquad \varphi\ \text{a negative literal}$

where the values of θ' and $\mathcal{D}'$ depend on occurrences of t and u:
- If $t \notin FV(\varphi[t/x], \Gamma_L) \cup \vec{l}$, then $\theta' := \theta$ and $\mathcal{D}' := \mathcal{D}$. In this subcase the precondition $t \notin FV(\varphi[t/x])$ implies that $x \notin FV(\varphi)$ and thus $\varphi[u/x] = \varphi[t/x]$.
- If $t, u \in FV(\varphi[t/x], \Gamma_L) \cup \vec{l}$, then $\theta' := \theta \wedge t = u$ and $\mathcal{D}' := \mathcal{D}$.
- Else it holds that $t \in FV(\varphi[t/x], \Gamma_L) \cup \vec{l}$ and $u \notin FV(\varphi[t/x], \Gamma_L) \cup \vec{l}$. Then $\theta' := \theta[t/u]$ and $\mathcal{D}' := \mathcal{D}[t/u]^{\mathsf{RHS}}$. To derive property I2. for this subcase, that is, $\models \neg\theta[t/u] \vee \varphi[t/x], \Gamma_L \vee \mathcal{D}[t/u]^{\mathsf{RHS}}$, it is required that $u \notin FV(\mathcal{D}[t/u]^{\mathsf{RHS}})$, which follows from the precondition $u \notin \vec{l}$ of the subcase.

This completes the proof of Theorem K.1.