Title
Geometric Model Theory in Efficient Computability

Author
Hill, Cameron Donnay

Publication Date
2010

Geometric Model Theory in Efficient Computability

by

Cameron Donnay Hill

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Leo Harrington, Chair
Professor Thomas Scanlon
Professor Luca Trevisan

Fall 2010

Geometric Model Theory in Efficient Computability

Abstract

Geometric Model Theory in Efficient Computability

by

Cameron Donnay Hill

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Leo Harrington, Chair

This dissertation consists of the proof of a single main result linking geometric ideas from the first-order model theory of infinite structures with complexity-theoretic analyses of problems over classes of finite structures. More precisely, we show that for a complete finite-variable theory of finite structures, models are efficiently recoverable from elementary diagrams if and only if the theory is super-rosy. In the course of the argument, we reconstitute the machinery of þ-independence and rosiness for classes of finite-structures, as well as a characterization of rosy classes analogous to the Independence theorem for the simple theories. We show that a super-rosy theory admits a weak form of model-theoretic coordinatization, which can be converted into to an algorithm for the model-building problem mentioned above in a natural and intuitive way. Conversely, we show how to extract a model-theoretic independence relation directly from an efficient algorithm for the model-building problem.

Mom

Contents

Acknowledgments

I thank my primary advisor, Leo Harrington, for helping me to find the kind of mathematician that I apparently am. I also thank him for helping me to realize that I don't need to throttle the "mundane" mathematician in me in order to be that fellow.

I thank my other advisor, Thomas Scanlon, for his help at very key moments in the proof of the conjecture, "I will write a thesis."

I thank Anuj Dawar and Bjarki Holm for tolerating my insane-sounding babbling and also for some very useful conversations.

I thank the staff of the UC Berkeley Mathematics department for their help over these years, especially Barb Waller, Marsha Snow and Barbara Peavy.

Chapter 1

Introduction

1.1 Some introductory remarks

There is a significant and steadily growing literature on what may be called, roughly, structural analysis of "easy" instances of hard computational problems, and there are, to my mind, two main currents in this study. The first of these – and also the more developed, it seems – is identifiable as a kind of algorithm design practice [13], wherein the analyst isolates some structural property and exploits this property to design an efficient algorithmic solution of the problem of interest when restricted to the corresponding sub-class of instances. A prime example of this impulse is found in "parametrized complexity theory," where the structural properties under scrutiny – bounded tree-width, branch-width, rank-width and so forth, and the associated decomposition templates – are commonly lifted from (finite) graph theory. These graph-based structural notions *are* applicable when computability is examined in the logical framework of finite model theory, but in general, they do not easily fit into the machinery of geometric model theory, which is my primary device here. The second current might be understood as a converse of the first; the analyst attempts to recover structure in a sub-class of problem instances from the hypothesis of efficient computability directly or from the efficacy of a certain style of algorithm. A central, if implicit, objective of both currents is to see the two meet, thereby giving an exact characterization of "easy" sub-classes of instances for a problem or of the domain of applicability of some style of algorithm.

There are at least two well-known successes of this last kind. In [5], we see that PAC-learnability[1] with a polynomially-bounded number of samples (in terms of the reciprocals of the accuracy parameters) is possible precisely when the concept class has finite Vapnik-Chervonenkis dimension. In [25] can be found a proof of the result – now part of the folklore of combinatorial optimization – that the efficacy of the greedy algorithm (naively making locally-optimal choices) corresponds exactly to the presence of matroid structure in the problem instances. Under restrictions of the nature of the objective function of the optimization

[1] probably approximately correct learning

problem, [18] and [14], for example, recover similar structural characterizations in terms of somewhat less ideal structures (greedoids and matroid embeddings, respectively). I personally find it suggestive that both of these conditions – finite VC-dimension and matroid structure – are already extensively studied in first-order model theory (of infinite structures). "Infinite VC-dimension" is known to model-theorists as the independence property, and "dependent" theories (theories which do not have the independence property) have come under close scrutiny in the last decade. Matroids, under the name "pregeometry," occur pervasively as the primitive or irreducible subsets of models in most (or all) model-theoretic structure theorems. The work of this dissertation, then, is an attempt to see these apparent connections between complexity theory and geometric model theory bear some fruit.

In fact, this dissertation essentially consists of the proof of a single main result linking geometric ideas from the first-order model theory of infinite structures with complexity-theoretic analyses of problems over classes of finite structures. To remove any suspense (and hopefully justify the effort of reading beyond the introduction), the statement of the theorem is as follows:

Theorem (8.24 of chapter 8)**.** *Let $K = \textbf{\textit{fin}}[T^G]$, where T is a complete k-variable theory with infinitely many finite models up to isomorphism.*

I. If T is constructible, then K is rosy.

II. T is efficiently constructible if and only if K is super-rosy.

Obviously, a great number of definitions are needed (regardless of the readers background, most likely) to make sense of these assertions. For the time being, it should be understood as a shadow of the "main current" of first-order model theory – namely, Shelah's Classification theory [27]. I take "efficiently constructible" – meaning that models of T can be efficiently recovered from elementary diagrams of subsets (chapter 4) – to be a reasonable substitute for "classifiable" in the classical theory. We then seek a hierarchy of structural properties culminating in efficient constructibility in analogy with the stability-theoretic hierarchy, Stable$\supsetneq$Super-stable$\supsetneq$Classifiable=Super-stable+NDOP. In the classical scenario, any non-trivial bound on the number of models of the theory in each cardinality imposes stability, which already supports the rudimentary notion of geometry known as non-forking independence. In the scenario of this study, the hypothesis of constructibility by an algorithm cursorily *imitating* that of an efficient algorithm in form (meaning, an essentially inflationary program which isn't necessarily efficient) is sufficient to impose another rudimentary notion of geometry on the class of models – in this case, known as þ-independence in a rosy class;[2] this is the content of *I* of the theorem. The further requirement of efficiency – polynomially-bounded running times – induces a further guarantee of good behavior in the geometry of þ-independence, and the "only if" portion of *II* of the theorem (and theorem 4.12) amounts to just this fact. It turns out, then, that this additional tractability in

[2]The character þ is pronounced "thorn."

the geometry gives enough purchase to devise an efficient algorithm, initially disguised as a weak model-theoretic coordinatization result (chapter 7), for the class of the theory's finite models.

1.2 Outline of the dissertation

Chapter 2: *Essentials of finite-variable logics.*

> We present the definition of the k-variable fragment of first-order logic and its infinitary variant, as well as the k-pebble game, which characterizes "elementary equivalence" and k-variable types in analogy with Ehrenfeucht-Fraïssé games for the full first-order logic. From the k-pebble game, we define the complete invariant for k-variable logic and its (more useful) relaxation, the theory T^G of game tableaux for a complete k-variable theory T. Finally, we define the notion of a *capped* k-variable theory, and prove that the class $\mathbf{fin}[T^G]$ of finite models of T^G has amalgamation over sets whenever T is capped.

Chapter 3: *Essentials of transformations of structures.*

> We define abstractly the concept of a transformation of structures over the class "hereditarily-finite lists," as well as the basic programming language of our analyses in succeeding chapters. With these definitions in place, and after mentioning some easy and/or folkloric normalizations of programs, we give our definition of what it means to "solve" the model-building problem for capped k-variable theory.
>
> A key idea throughout is that the algorithms in question are functorial and, more, functorial "all the way down" – meaning that all sub-routines and even the most primitive instructions of the programming language are themselves functorial.

Chapter 4: *Deeper characterizations of efficient transformations*

> We sketch a proof that an efficient solution of the model-problem amounts to an "essentially inflationary" solution. (We give only a sketch because the complete demonstration is not interestingly different from the analyses in [1].)
>
> We also argue that an *efficient* coherent solution for the model-building problem implies a certain property we call *small algebraicity*, which enters into our later analyses of super-rosy classes and the recovery of an efficient model-building algorithm for such classes.

Chapter 5: *þ-Independence and rosy Fraïssé classes.*

We introduce the notions of strong dividing, þ-dividing and þ-forking for Fraïse classes K of finite structures, as well as the local þ-rank. We define rosiness in terms of the local þ-rank. We show that local þ-rank characterizes þ-forking, and we use this fact to prove both that þ-independence is a weak notion of independence in every Fraïssé class and that it is a true independence relation if and only if K is rosy. Our analysis largely avoids the use of the Compactness theorem, but except in some technical details, the development differs little from the classical treatment in [22].

Chapter 6: *A further characterization of rosiness and examples.*

We give an additional key characterization of rosiness for Frašse classes – showing that rosiness is the coarsest possible notion of independence and that the presence of any true independence relation implies that the class is rosy. We also give definitions of *local character* and super-rosiness, and we analyze their relationship with another property which we call *small algebraicity.* We also include a sketch of some properties of the $U^{þ}$-rank for super-rosy Fraïssé classes.

Subsequently, we present two examples of rosy classes – finite vector spaces over a finite field and models of the k-variable fragment of the theory of the random graph. Finally, we sketch an argument showing that the theory of a random pseudo-scale [12], even though it is capped and has small algebraicity, is *not rosy* – thus demonstrating that the property of rosiness has some nontrivial content.

Chapter 8: *Coordinatization and efficient model-building*

Using techniques [Alf+Usvyatsov], we demonstrate a weak notion of coordinatization for a super-rosy Fraïssé class (equivalently, for a countably categorical super-rosy theory with finite $U^{þ}$-rank and which eliminates imaginaries). We use this notion of coordinatization to prove the converse of the main result of chapter 7 – that any capped super-rosy k-variable theory admits efficient model-building from elementary diagrams.

Chapter 7: *Unfolding digraphs of transformations and independence*

We define the naive unfolding graph of an essentially inflationary program acting on a structure, and we show how to prune away portions of it that obscure its model-theoretically interesting content. We then define the notions of d-separation and deviation of extensions of types, and we show that non-deviation is a notion of independence. Finally, we show that under conditions satisfied by a solution of the model-building problem for some capped k-variable theory T, this notion of independence is a true independence relation on $\mathbf{fin}[T^G]$. We conclude from this that if T admits *essentially inflationary* model-building from elementary diagrams, then $\mathbf{fin}[T^G]$ is a rosy class.

1.3 Notation and terminology

- **Logic notion**:

  Our notion for structures (in the model-theoretic sense) is fairly standard.

  - A signature ρ is a set of relation-symbols and constant symbols (usually none of the latter) with distinguished arities; obviously, constant symbols have arity 0. We use the notations $R^{(n)}$ and $ari(R) = n$ interchangeably to mean that R is a relation symbol arity n, and the notation $R^{(n)} \in \rho$ means that R is a relation symbol *of* ρ of arity n.
  - A ρ-structure, then, is a tuple $\mathcal{A} = (A, (R^{\mathcal{A}})_{R\in\rho}, (c^{\mathcal{A}})_{c\in\rho})$ where A is a set. $R^{\mathcal{A}} \subseteq A^n$ whenever $R^{(n)} \in \rho$, and $c^{\mathcal{A}} \in A$ whenever $c \in \rho$ is a constant symbol.

    Usually, a structure is denoted by a script capitol letter – like $\mathcal{A}$ – and its universe, A, is understood from context. In case the structure's universe is not clear, we write $||\mathcal{A}||$ to denote the universe of a structure $\mathcal{A}$. If we need to consider the cardinality of the universe of $\mathcal{A}$ in this scenario, we write $\#||\mathcal{A}||$ instead of $|||\mathcal{A}|||$
  - $\mathbf{fin}[\rho]$ denotes the class of finite ρ-structures, and if T is a theory, then $\mathbf{fin}[T]$ denotes the class of finite models of T.
  - If $\mathcal{M}$ is a structure and $A \subseteq M$, then we write $(A; \mathcal{M})$ as an abbreviation of $diag^{\mathcal{M}}(A)$ – the quantifier-diagram of A with respect to $\mathcal{M}$.

- Sets and tuples

  - For a positive finite number n, we set $[n] = \{1, ..., n\}$
  - A tuple $\overline{a}$ over a set A is understood to be a function $\overline{a} : [n] \to A$ for some positive finite number n. Thus, A^n is formally identical to ${}^{[n]}A$, the set of functions $[n] \to A$, and if $\overline{a} \in A^n$, $rng(\overline{a}) = \{\overline{a}(i) : i \in [n]\}$ is a subset of A. For convenience, we also write $\overline{a} = (a_1, ..., a_n)$ with the understanding that $\overline{a}(i) = a_i$ for each $i \in [n]$.
  - If X is a set and n is a positive number, $\binom{X}{n}$ and $\binom{X}{\leq n}$ denote the families of subsets of cardinality exactly n and cardinality at most n, respectively.

    Note that $\binom{[m]}{n}$ is the set of subsets of $[m] = \{1, ..., m\}$ of size exactly n, and the cardinality of $\binom{[m]}{n}$ is $\binom{m}{n}$
  - If M is a (possibly infinite) set, we write $X \subset_{\mathbf{fin}} M$ to mean that X is a *finite* subset of M. $\mathscr{P}_{\mathbf{fin}}(M)$ denotes the set of all finite subsets of M, and

    $$M^{<\omega} = \{\, ()\,\} \bigcup_{0<n<\omega} M^n$$

    where $()$ is the "empty tuple," which *we do not* identify with the empty set.

  - As is fairly standard in the model theory literature, we write AB as shorthand for $A \cup B$.
  - We write $|X|$ for the cardinality of a set X (which will almost always be a finite number) except in the situation mentioned above.
  - Suppose $0 < k < l < \omega$ and $\overline{a} \in A^k$, say $\overline{a} = (a_1, ..., a_k)$. Then we define

$$pad_l(\overline{a}) = (a_1, ..., a_k, \underbrace{a_k, ..., a_k}_{l-k \text{ times}})$$

- Big-O, Big-Ω and Big-Θ notaion.

  We make relatively scant use of these notations, but the interpretation of the expression $\Omega(g(n))$ seems to be somewhat ambiguous. For functions $f, g : \omega \to \omega$, we assert the following definitions

  - $f(n) \in O(g(n))$ if there are $n_0 < \omega$ and $\delta \in \mathbb{R}$ such that $f(n) \leq \delta g(n)$ whenever $n_0 \leq n < \omega$.
  - $f(n) \in \Omega(g(n))$ if there are $n_0 < \omega$ and $\delta \in \mathbb{R}$ such that $f(n) \geq \delta g(n)$ whenever $n_0 \leq n < \omega$.
  - $f(n) \in \Theta(g(n))$ if both $f(n) \in O(g(n))$ and $f(n) \in \Omega(g(n))$.

- Independence

  There is a small difficulty in giving a precise definition of either a "notion of independence" or a "true independence relation" before exposing the material chapter [not this one]. Specifically, we have not yet asserted the meaning of a type. For the moment, we will gloss the meaning of “equivalence” in the statements of Invariance and Extension below.

  Be that as it may, the intuition behind the terminology is as follows. A *notion of independence* corresponds to a proposal for a notion of geometry in a class of structures, and the axioms associated with the term are those that should be “easy” to satisfy and verify. Moreover, we expect that these are, then, properties of the notion itself, requiring nothing further from the class under examination. A notion of independence may then be a *true independence relation* – a useful concept of geometry in the class – when certain other properties are satisfied by the class. For example, in chapter 5, we find that þ-independence, $\mathop{\smile\!\!\!\!\!\!\mid}^{þ}$, is always a notion of independence, as we'd expect from the terminology, and $\mathop{\smile\!\!\!\!\!\!\mid}^{þ}$ is a true independence relation just in case the class is rosy.

  Of course, these axioms (along with the Local character axiom) first appeared in the articles leading up to [27]. To my knowledge, [17, 16] together form the only previous treatment of model-theoretic geometry for classes of finite structures; these inform some of the setup for our analysis but play little or no direct role. My choice of (and names of) the axioms derives [28, 22, 10, 2].

- Notion of independence, $\mathop{\smile\hskip-0.9em ^| \ }^{\circ}$:
    1. Invariance. If $(A, B, C) \equiv (A_1, B_1, C_1)$ and $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B$, then $A_1 \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C_1} B_1$.
    2. Extension.
       If $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B$ and $BC \subseteq D$, then there are $\mathcal{M}' \in K_D$ and $A' \subseteq M'$ such that $A' \equiv_{BC} A$ and $A' \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} D$.
    3. Monotonicity. If $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B$ and $B_0 \subseteq B$, then $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B_0$.
    4. Base-monotonicity. If $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B$ and $B_0 \subseteq B$, then $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{CB_0} B$.
    5. Partial right-transitivity. If $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B_1 B_2$, then $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B_1$ and $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{CB_1} B_2$.
    6. Preservation of algebraic dependence I. If $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} A$, then $A \subseteq acl(C)$
       Preservation of algebraic dependence II. If $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B$, then $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} acl(B)$.
       Preservation of algebraic dependence III.
       If $B \cap acl(AC) \setminus acl(A)$ is non-empty, then $A \not\mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B$
- True independence relation.
    1. Axioms of a notion of independence.
    2. Existence. $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} C$.
    3. Symmetry. If $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B$, then $B \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} A$
    4. Full transitivity. $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B_1 B_2$ if and only if $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{C} B_1$ and $A \mathop{\smile\hskip-0.9em ^| \ }^{\circ}_{CB_1} B_2$.

Chapter 2

Essentials of finite-variable logics

One of our concessions in return for a gentler analysis is the move to considering finite-variable-elementary classes. To some degree, this move is similar to the "decision" in classical model theory to develop a machinery of classification theory for complete first-order theories before tackling non-elementary classes of structures. In section 1 of this chapter, then, we summarize the basic facts about finite-variable logics; in section 2 and thereafter, we extend this machinery a little for a slightly restricted subset of complete finite-variable theories.

2.1 The very basics of finite-variable logics

2.1.1 The logics L^k and $L^k_{\infty,\omega}$, and the k-pebble game

As will almost always be the case, let ρ be a finite signature with no function symbols; let C_ρ denote the set of constant symbols of ρ. Further assume that $k < \omega$ is not less that $ari(R)$ for every relation symbol R of ρ, and in any case, $k \geq 2$. Let $V = \{x_1, ..., x_k\}$ be a set of variables. The k-variable fragment of first-order logic, $L^k = L(\rho)^k$, is defined to be the smallest set S of formulas satisfying the following:

1. If $s, t \in V \cup C_\rho$, then $(s = t) \in S$.

2. If $R^{(n)} \in \rho$, $n \leq k$, and $s : n \to V \cup C_\rho$, then $R(s(0), ..., s(n-1)) \in S$.

3. If $\varphi_1, \varphi_2 \in S$ and $x \in V$, then $(\varphi_1 \wedge \varphi_2)$, $(\neg\varphi_1)$ and $(\exists x\, \varphi_1)$ are in S.

The infinitary k-variable logic is obtained in the same manner with the addition of a fourth inductive formation rule:

4. If α is an ordinal and $\varphi_i \in L^k$ for each $i < \alpha$, then $\bigwedge_{i<\alpha} \varphi_i$ and $\bigvee_{i<\alpha} \varphi_i$ are in S

This set of (infinitary) formulas is, then, denoted $L^k_{\infty,\omega} = L(\rho)^k_{\infty,\omega}$.[1]

In analogy with the Ehrenfeucht-Fraïssé game in first-order logic, the *k-pebble game* characterizes both L^k-elementary equivalence, $\equiv^k$, and (L^k, k)-types with respect to T.[2] We now describe the k-pebble game. There are two players, the Spoiler and the Duplicator, and they play on a pair of (not necessarily distinct) ρ-structures $\mathcal{A}$ and $\mathcal{B}$. We also maintain two *partial* functions $f_A : [k] \rightharpoonup A$ and $f_B : [k] \rightharpoonup B$ – call them pebble functions – which, at the outset, are empty. The k-pebble game $PG_k(\mathcal{A}, \mathcal{B})$ consists of the following protocol. Play proceeds in rounds, and each round has the following form:

1. The Spoiler selects a structure – for convenience, say he chooses $\mathcal{A}$. He then chooses a number $i \in [k]$ and an element $a \in A$. We then modify f_A by (re)setting $f_A(i) := a$.

2. In response, the Duplicator chooses an element of the other structure – in this case, some $b \in B$ – and we (re)set $f_B(i) := b$

Note that at the end of each round, the domains of f_A and f_B are the same, say $D = dom(f_A) = dom(f_B)$. The round is a win for the Duplicator (she survives the round) just in case the map $g = \{(f_A(i), f_B(i)) : i \in D\}$ is a partial isomorphism (i.e. it is a bijection between $rng(f_A)$ and $rng(f_B)$, and for every atomic formula $\varphi(\overline{x})$ and $\overline{a} \in rng(f_A)^{\overline{x}}$, $\mathcal{A} \vDash \varphi(\overline{a})$ iff $\mathcal{B} \vDash \varphi(g\overline{a})$). Temporarily, we write $\mathcal{A} \sim^k \mathcal{B}$ if the Duplicator has a strategy which permits her to carry on the game indefinitely (for ω rounds) such that she wins each round – in brief, the Duplicator has a winning strategy in $PG_k(\mathcal{A}, \mathcal{B})$ for ω rounds. Naturally, the game can begin with pebble functions pre-loaded with elements of the corresponding structures. Thus, if $l \leq k$, $\overline{a} \in A^l$ and $\overline{b} \in B^l$, we write $PG_k(\mathcal{A}, \overline{a}; \mathcal{B}, \overline{b})$ for the game which begins with $f_A = \{(i, a_i) : i \in [l]\}$ and $f_B = \{(i, b_i) : i \in [l]\}$. Assuming the Duplicator has a winning strategy in $PG_k(\mathcal{A}, \overline{a}; \mathcal{B}, \overline{b})$, we write $(\mathcal{A}, \overline{a}) \sim^k (\mathcal{B}, \overline{b})$, or just $\overline{a} \sim^k \overline{b}$ when $\mathcal{A} = \mathcal{B}$. We collect the results regarding the k-pebble game in the following theorem:

Theorem 2.1. *Let $\mathcal{A}$ and $\mathcal{B}$ be finite structures.*

1. *$\mathcal{A} \sim^k \mathcal{B}$ if and only if $\mathcal{A}$ and $\mathcal{B}$ agree on every sentence of L^k if and only if $\mathcal{A}$ and $\mathcal{B}$ agree on every sentence of $L^k_{\infty,\omega}$.*

   (The latter two conditions are denoted $\mathcal{A} \equiv^k \mathcal{B}$.)

2. *Suppose $l \leq k$, $\overline{a} \in A^l$ and $\overline{b} \in B^l$. Then $(\mathcal{A}, \overline{a}) \sim^k (\mathcal{B}, \overline{b})$ if and only if $tp^k(\overline{a}; \mathcal{A}) = tp^k(\overline{b}; \mathcal{B})$.*

(Proofs of theorem 2.1 can be found in a number of texts, including [19] and [24].)

[1] In almost all of our work, no ambiguity arises when we write L^k, $L^k_{\infty,\omega}$ and dispense with indicating the signature.

[2] $tp^k(a_1, ..., a_k; \mathcal{A})$ – the (L^k, k)-type of $\overline{a} = (a_1, ..., a_k)$ in $\mathcal{A}$ – is just the set of L^k-formulas satisfied by $\overline{a}$ in $\mathcal{A}$. We will deal with types of higher arity and types over sets of parameters later in this chapter.

The remainder of this subsection is not hugely important for the main thread of the dissertation, but the reader may find it mildly interesting. Over finite structures, $L^k_{\infty,\omega}$-formulas have an especially simple normal form:

Proposition 2.2 (see [19]). *For every formula $\varphi(\overline{x})$ of $L^k_{\infty,\omega}$, there are countable families $\{\varphi_i(\overline{x})\}_{i<\omega}$, $\{\varphi'_i(\overline{x})\}_{i<\omega}$ of L^k-formulas such that for any finite structure $\mathcal{A}$,*

$$\mathcal{A} \vDash \varphi(\overline{a}) \Leftrightarrow \mathcal{A} \vDash \bigvee_{i<\omega} \varphi_i(\overline{a}) \Leftrightarrow \mathcal{A} \vDash \bigwedge_{i<\omega} \varphi'_i(\overline{a})$$

for all $\overline{a} \in A^{\overline{x}}$.

Suppose $\Theta = \{\theta_i(\overline{x}_i)\}_{i<m}$ is a family of $L^k_{\infty,\omega}$ formulas. $\mathcal{A}[\Theta]$ denotes the expansion of $\mathcal{A}$ in a signature $\rho[\Theta] = \rho \dot{\cup} \{R_i : i < m\}$ such that, naturally enough, $R_i^{\mathcal{A}[\theta]} = \{\overline{a} \in A^{\overline{x}_i} : \mathcal{A} \vDash \theta_i(\overline{a})\}$ for each $i < m$. If $m = 1$, so that $\Theta = \{\theta(\overline{x})\}$, then we also write $\mathcal{A}[\theta]$ in place of $\mathcal{A}[\Theta]$.

Lemma 2.3. *Let $\theta(\overline{x}) \in L^k_{\infty,\omega}$ be a conjunction of countably many L^k-formulas. Then for all ρ-structures $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{A} \equiv^k \mathcal{B}$ implies $\mathcal{A}[\theta] \equiv^k \mathcal{B}[\theta]$.*

Proof. Suppose σ is a strategy for the Duplicator witnessing $\mathcal{A} \equiv^k \mathcal{B}$. We claim that σ also witnesses $\mathcal{A}[\theta] \equiv^k \mathcal{B}[\theta]$. For convenience, we assume that θ is k-ary, but this is inessential. Consider the end of a round in which (w.l.o.g.) $dom(f_A) = dom(f_B) = k$. Since σ allows the Duplicator to survive ω rounds from this position, it follows that

$$(\mathcal{A}, f_A(1), ..., f_A(k)) \equiv^k (\mathcal{B}, f_B(1), ..., f_B(k))$$

so $tp^k(f_A(1), ..., f_A(k); \mathcal{A}) = tp^k(f_B(1), ..., f_B(k); \mathcal{B})$. In particular,

$$\mathcal{A} \vDash \theta(f_A(1), ..., f_A(k)) \Leftrightarrow \mathcal{B} \vDash \theta(f_B(1), ..., f_B(k))$$

and this proves the claim. □

As an easy corollary, we have:

Proposition 2.4. *Consider a family $\Theta = \{\theta_i(\overline{x}_i)\}_{i<m}$ of $L^k_{\infty,\omega}$-formulas. Then for all structures $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{A} \equiv^k \mathcal{B}$ implies $\mathcal{A}[\Theta] \equiv^k \mathcal{B}[\Theta]$.*

2.1.2 Types and L^k-elementary embeddings

While k-variable n-types for $n \leq k$ are straightforward to define, for $n > k$, we need to be slightly more careful. (We follow [17] in our definition.) Let $0 < n < \omega$, let $v_1, ..., v_n$ be a sequence of pairwise distinct variables. Let $V = \{x_1, ..., x_k\}$ be the set of variables used in the construction of L^k. An *n-ary augmented L^k-formula* is a pair $(\varphi(x_1, ..., x_k), f)$ where

$\varphi \in L^k$ and $f : [k] \to [n]$; we write $[\varphi, f](v_1, ..., v_n)$ for this object. If $\mathcal{A}$ is a ρ-structure and $a_1, ..., a_n \in A$, then we define

$$\mathcal{A} \vDash [\varphi, f](a_1, ..., a_n) \;\Leftrightarrow\; \mathcal{A} \vDash \varphi(a_{f(1)}, ..., a_{f(k)})$$

Let $aug_n L^k$ denote the set of all n-ary augmented L^k-formulas. The complete L^k-type of $\overline{a}$ over $\emptyset$ with respect to $\mathcal{A}$ – denoted $tp^k(\overline{a}; \mathcal{A})$ or $tp^k(\overline{a})$ if $\mathcal{A}$ is clear from context – is the following set of n-ary augmented L^k-formulas:

$$\left\{[\varphi, f](v_1, ..., v_n) \in aug_n L^k : \mathcal{A} \vDash [\varphi, f](\overline{a})\right\}$$

Similarly, if $C \subseteq A$, then the complete L^k-type of $\overline{a}$ over C with respect to $\mathcal{A}$ – denoted $tp^k(\overline{a}/C; \mathcal{A})$ or $tp^k(\overline{a}/C)$ – is the set

$$\bigcup_{m<\omega} \bigcup_{\overline{c} \in C^m} \left\{[\varphi, f](\overline{v}, \overline{c}) : [\varphi, f] \in aug_{n+m} L^k,\, \mathcal{A} \vDash [\varphi, f](\overline{a}, \overline{c})\right\}$$

Augmented L^k-formulas with parameters of the form of those occurring in $tp^k(\overline{a}/C)$ are called *n-ary augmented formulas over C*. The following observation is quite useful but more or less self-evident.

Observation. *Let $\mathcal{A}$ be a ρ-structure, $C \subseteq A$ and $\overline{a}, \overline{b} \in A^n$. The following are equivalent:*

1. $tp^k(\overline{a}/C; \mathcal{A}) = tp^k(\overline{b}/C; \mathcal{A})$
2. *For all $l \leq k$, $i_1, ..., i_l \in [n]$ and $c_{l+1}, ..., c_k \in C$,*

$$tp^k(a_{i_1}, ..., a_{i_l}, c_{l+1}, ..., c_k; \mathcal{A}) = tp^k(b_{i_1}, ..., b_{i_l}, c_{l+1}, ..., c_k; \mathcal{A})$$

Suppose K is a class of finite ρ-structures. A *K-realizable (L^k, n)-type over C with respect to $\mathcal{A}$* (where $C \subseteq A$ and $\mathcal{A} \in K$) is a set $\pi = \pi(v_1, ..., v_n)$ of n-ary augmented formulas over C such that there is a structure $\mathcal{B} \in K$ and an n-tuple $\overline{b} \in B^n$ such that

- For all $\overline{c} \in C^k$ and all $\varphi(x_1, ..., x_k) \in L^k$, $\mathcal{A} \vDash \varphi(\overline{c})$ if and only if $\mathcal{B} \vDash \varphi(\overline{c})$. (We say that $\mathcal{B}$ preserves $(C; \mathcal{A})$.)
- For all $[\varphi, f](\overline{v}, \overline{c}) \in \pi$, $\mathcal{B} \vDash [\varphi, f](\overline{b}, \overline{c})$.

In this case, we say that $\overline{b}$ realizes π in $\mathcal{B}$, and of course, it is not difficult to see that $\pi \subseteq tp^k(\overline{b}/C; \mathcal{B})$. If π is an (L^k, n)-type over C with respect to $\mathcal{A}$, we say that π is *maximal* if for each n-ary augmented formula over C, say $[\varphi, f](\overline{v}, \overline{c})$, either $[\varphi, f](\overline{v}, \overline{c}) \in \pi$ or $[\neg\varphi, f](\overline{v}, \overline{c}) \in \pi$; obviously, if π is K-realizable, then it's not the case that $[\varphi, f](\overline{v}, \overline{c}) \in \pi$ and $[\neg\varphi, f](\overline{v}, \overline{c}) \in \pi$. It's relatively easy to see that if all models under consideration are finite, then a maximal K-realizable (L^k, n)-type $\pi(\overline{v})$ over C with respect to $\mathcal{A}$ is always the complete L^k-type of some n-tuple $\overline{b} \in B^n$ for some structure $\mathcal{B} \in K$ which preserves $(C; \mathcal{A})$.

A note for future reference: We will often use the notation $K_{(C;\mathcal{A})}$, and more briefly K_C, to denote the sub-class of structures $\mathcal{B} \in K$ such that $C \subseteq B$ and $\mathcal{B}$ preserves $(C; \mathcal{A})$.

Suppose $\mathcal{M}$ and $\mathcal{N}$ are finite ρ-structures, and let $A \subseteq M$ and $h : A \to N$. We say that h is a *partial L^k-elementary embedding* if for every $\overline{a} \in A^k$ and every $\varphi(x_1, ..., x_k) \in L^k$, $\mathcal{M} \models \varphi(\overline{a})$ if and only if $\mathcal{N} \models \varphi(h\overline{a})$. Such a map h is an L^k-elementary embedding just in case $A = M$.

Observation. *Suppose $\mathcal{M}$ and $\mathcal{N}$ are finite ρ-structures, and let $A \subseteq M$ and $h : A \to N$. Let $\overline{a}$ be an enumeration of A. Then h is a partial L^k-elementary embedding if and only if $tp^k(\overline{a}; \mathcal{M}) = tp^k(h\overline{a}; \mathcal{N})$.*

2.2 The complete invariant for L^k and game tableaux

In this subsection, we introduce both the complete invariant I^k for k-variable theories of finite ρ-structures and the notion of a game tableau for the L^k-theory of a fixed finite ρ-structure. The latter is key in much of our analysis of computational problems around finite-variable logic. The invariant is not itself terribly useful in this work, but as well as providing the starting point for game tableaux, it is foundational to the L^k-canonization problem, which is ancestral to our model-building problem. The material on the complete invariant can be found in [19] or [24], and the material on game-tableaux can be found (with non-trivial differences) in [24], where it is deployed only in relation to the 2-variable logic.

Suppose $\mathcal{M}$ is a finite ρ-structure. By theorem 2.1, the quotient structure $M^k/{\equiv^k}$ is essentially synonymous with the set $S^k_k(T)$ of k-variable k-types of T, where $T = Th^k(\mathcal{M})$ is the complete k-variable theory of $\mathcal{M}$. Moreover, if the set of quantifer-free k-types of ρ is endowed (*a priori* but arbitrarily) with a linear order, then there is an algorithm computing a mapping

$$\mathbf{fin}[\rho] \longrightarrow \mathbf{fin}[\{<\}] : \mathcal{N} \mapsto (N^k/{\equiv^k}, <^{\mathcal{N}})$$

where $(N^k/{\equiv^k}, <^{\mathcal{N}})$ is a linear order;[3] this algorithm has running-time no worse than $O(|N|^{3k})$ (see [24] or [19]). Now, for each complete quantifier-free k-type $\theta(x_1, ..., x_k)$ of ρ, let $V^{(1)}_\theta$ be a new unary predicate symbol; for each permutation $\sigma \in Sym[k]$, let $P^{(2)}_\sigma$ be a new binary predicate symbol; and let $Acc^{(2)}$ be an additional unary predicate symbol. Let ρ^{inv} be the signature consisting of these symbols together with the binary relation symbol $<$. (Note that $\rho \nsubseteq \rho^{\text{inv}}$ and $\rho^{\text{inv}} \nsubseteq \rho$.) Given a finite ρ-structure $\mathcal{M}$, then, we define $I^k(\mathcal{M})$ to be a ρ^{inv}-structure with universe $M^k/\equiv^k$ as follows:

- $<^{I^k(\mathcal{M})}$ is the linear order, $<^{\mathcal{M}}$, of $M^k/{\equiv^k}$ described above.
- $V_\theta^{I^k(\mathcal{M})} = \left\{\overline{a}/{\equiv^k} : \overline{a} \in M^k,\ \mathcal{M} \models \theta(\overline{a})\right\}$ for each quantifier-free k-type θ.

[3] Here and in the sequel, $\mathbf{fin}[\rho]$ denotes the class of finite ρ-structures.

- For $\sigma \in Sym[k]$, we put $(\overline{a}/\equiv^k, \overline{b}/\equiv^k) \in P_\sigma^{I^k(\mathcal{M})}$ just in case,

$$(\mathcal{M}, (a_{\sigma(1)}, ..., a_{\sigma(k)})) \equiv^k (\mathcal{M}, \overline{b})$$

- For $\overline{a}, \overline{b} \in M^k$, we put $(\overline{a}/\equiv^k, \overline{b}/\equiv^k) \in Acc^{I^k(\mathcal{M})}$ just in case there is an element $m \in M$, such that

$$(\mathcal{M}, (m, a_2, ..., a_k)) \equiv^k (\mathcal{M}, \overline{b})$$

  Equivalently, $(\overline{a}/\equiv^k, \overline{b}/\equiv^k) \in Acc^{I^k(\mathcal{M})}$ if for every $(a'_1, ..., a'_k) \in \overline{a}/\equiv^k$, there is an $m' \in M$ such that

$$(\mathcal{M}, (m', a'_2, ..., a'_k)) \equiv^k (\mathcal{M}, \overline{b})$$

This operator $I^k(-)$ is known in the literature as the *complete invariant* for k-variable logic.

Theorem 2.5. *([24]) Let $\mathcal{M}$ and $\mathcal{N}$ be finite ρ-structures. Then $\mathcal{M} \equiv^k \mathcal{N}$ if and only if $I^k(\mathcal{M}) \cong I^k(\mathcal{N})$. Moreover, the mapping $\mathcal{M} \mapsto I^k(\mathcal{M})$ is computable with running-time $|M|^{O(1)}$.*

We note that since $I^k(\mathcal{M})$ is linearly ordered, there is a canonically isomorphic ρ^{inv}-structure with universe $[n]$, where $n = |I^k(\mathcal{M})| < \omega$, in which the linear order is the standard one, and this transformation is computable in polynomial-time. Thus, it is not terribly abusive to write $I^k(\mathcal{M}) = I^k(\mathcal{N})$ instead of $I^k(\mathcal{M}) \cong I^k(\mathcal{N})$, with the understanding that we pass to this canonical or standard model.

Digression*: L^k-canonization.* Having described the L^k-invariant, we are, at last, in a position to describe the problem which initiated the work in this dissertation – the so-called L^k-canonization problem. For background, we begin with a more standard complexity-theoretic question. The hardness of the Graph-Isomorphism problem, stated below, is a major open question in complexity theory:

Graph-Isomorphism: *Given* a pair of finite graphs G_1 and G_2, *accept* iff $G_1 \cong G_2$.

Graph-Isomorphism is a natural candidate for an NPtime-problem which is neither in Ptime nor NPtime-complete; such a problem must exist if Ptime $\neq$ NPtime. A related computational problem, which would certainly provide a Ptime-solution for Graph-Isomorphism, lies in defining a Ptime-computable operator F from finite graphs to finite graphs such that:

- $F(G) \cong G$, and if $|V(G)| = n$, then $V(F(G)) = [n]$.
- Isomorphism invariance: If $G_1 \cong G_2$, then $F(G_1) = F(G_2)$.

$F(G)$ can be thought of as a canonical encoding of the isomorphism-type of G. Clearly, the existence of an efficient graph-canonization operator immediately shows that Graph-Isomorphism is in PTIME, and consequently, studying the possibility of defining such an operator is little easier than settling the status of Graph-Isomorphism itself.

However, graph-canonization does lend itself to interesting relaxations. One of these is the so-called L^k-canonization problem, wherein one seeks to define an efficient operator $F : \mathbf{fin}[\rho] \to \mathbf{fin}[\rho]$ such that

- $F(\mathcal{M}) \equiv^k \mathcal{M}$, and $||F(\mathcal{M})|| = [n]$ for some $n \in \mathbb{N}$ (possibly $n \neq |M|$).
- $\equiv^k$-invariance: If $\mathcal{M}_1 \equiv^k \mathcal{M}_2$, then $F(\mathcal{M}_1) = F(\mathcal{M}_2)$

Thus, $F(\mathcal{M})$ can be thought of as a *canonical model* of $Th^k(\mathcal{M})$. Of course, the invariant I^k already meets the second criterion of an L^k-canonization operator – $I^k(\mathcal{M}_1) = I^k(\mathcal{M}_2)$ if and only if $\mathcal{M}_1 \equiv^k \mathcal{M}_2$. It would seem that L^k-canonization is, then, a "simple matter" of inverting the map $\mathcal{M} \mapsto I^k(\mathcal{M})$; that is, we need only seek an efficient operator $H^k : \mathbf{fin}[\rho^{\text{inv}}] \rightharpoonup \mathbf{fin}[\rho]$ such that $I^k(H^k(I^k(\mathcal{M}))) = I^k(\mathcal{M})$. Perhaps unsurprisingly, this is, in fact, no simple matter at all:

> [24] I^2 is efficiently invertible.
>
> [12] For $k \geq 3$, there is no recursive inverse H^k of I^k.
>
> [20] For any $k \geq 2$, I^k is recursively invertible when restricted to the class of finite ρ-structures $\mathcal{M}$ such that $Th^k(\mathcal{M})$ is stable and has the amalgamation property.
>
> [21] For any $k \geq 2$, I^k is recursively invertible when restricted to the class of finite ρ-structures $\mathcal{M}$ such that $Th^k(\mathcal{M})$ is simple with trivial forking-dependence and has the amalgamation property.

To cut this story short, the test problem for our analyses, which we call the model-building problem, has the problem of inverting I^k as a special case in the sense that computing $H^k(I^k(\mathcal{M}))$ amounts to computing a model of $Th^k(\mathcal{M})$ from the empty induced substructure. (We also grant each of our algorithms a great deal more specific information about the *single* theory for which it is designed.) The model-building problem, and solutions of it, for a given theory are described in detail in chapters 2 and 3. $\square$

For many of the analyses in this dissertation, the complete invariant carries much unnecessary information, which has the tendency to obscure what is really essential. Thus, we reduce the complete invariant to a less informative structure called a *game-tableau theory*. As we have noted before, the quotient set $M^k/{\equiv^k}$ is essentially synonymous with the set $S^k_k(T)$ of k-variable k-types of the theory $T = Th^k(\mathcal{M})$. Moreover, the *accessibility relation* between k-variable k-types of $\mathcal{M}$ is an invariant of T. That is, we may consider

$Acc \subseteq S_k^k(T) \times S_k^k(T)$ such that for any ρ-structure $\mathcal{N}$, if $\mathcal{N} \vDash T$, $\bar{b} \in N^k$, $p = tp^k(\bar{b};\mathcal{N})$ and $(p,q) \in Acc$, then there is an $n \in N$, such that $tp^k(n, b_2, ..., b_k; \mathcal{N}) = q$, and moreover, if $n' \in N$ and $q' = tp^k(n', b_2, ..., b_k; \mathcal{N})$, then $(p, q') \in Acc$. Similarly, if $\sigma \in Sym[k]$ and $tp^k(\overline{a};\mathcal{M}) = tp^k(\bar{b};\mathcal{N})$, then

$$tp^k(a_{\sigma(1)}, ..., a_{\sigma(k)}; \mathcal{M}) = tp^k(b_{\sigma(1)}, ..., b_{\sigma(k)}; \mathcal{N})$$

In fact, as we shall shortly see, these facts together with some types in the language of equality effectively determine the class of models of the theory T.

Let $\mathcal{M}_0$ be a fixed finite ρ-structure, and let $T = Th^k(\mathcal{M}_0)$. We enumerate $S_k^k(T)$ (arbitrarily) by $\alpha_1(\overline{x}), ..., \alpha_N(\overline{x})$, where $\overline{x} = (x_1, ..., x_k)$ is a tuple of pairwise distinct variables (say, an enumeration of the set V of variables used in the construction of the logic L^k). For each $\alpha \in S_k^k(T)$, let $R_\alpha^{(k)}$ be a k-ary relation symbol, and let $\mu_\alpha(\overline{x})$ be the unique complete (quantifier-free) k-type in the language of equality such that $T \vDash \forall \overline{x}(\alpha(\overline{x}) \to \mu_\alpha(\overline{x}))$. Let $\rho^G = \left\{R_\alpha : \alpha \in S_k^k(T)\right\}$. We define T^G to be the theory in the language of ρ^G consisting of the following assertions:

G1: $\forall x_1 ... x_k \bigvee_\alpha \left(R_\alpha(\overline{x}) \wedge \neg \bigvee_{\beta \neq \alpha} R_\beta(\overline{x})\right)$

G2: The "type" R_α of a k-tuple matches the equality type of the genuine type α: $\bigwedge_\alpha \forall x_1 ... x_k \left(R_\alpha(\overline{x}) \to \mu_\alpha(\overline{x})\right)$

G3: $\bigwedge_{\sigma \in Sym[k]} \bigwedge_\alpha \forall x_1 ... x_k \left(R_\alpha(\overline{x}) \leftrightarrow R_{\alpha^\sigma}(x_{\sigma(1)}, ..., x_{\sigma(k)})\right)$.

We write α^σ for the unique type β such that

$$T \vDash \forall x_1 ... x_k \left(\alpha(\overline{x}) \leftrightarrow \beta(x_{\sigma(1)}, ..., x_{\sigma(k)})\right)$$

G4: $\forall x_1 ... x_k \forall y \left(R_\alpha(\overline{x}) \to \bigvee_{\beta \in Acc(\alpha, -)} R_\beta(y, x_2, ..., x_k)\right)$

G5: $\bigwedge_\alpha \exists x_1 ... x_k \left(R_\alpha(\overline{x})\right)$

G6: $\bigwedge_\alpha \bigwedge_{\beta \in Acc(\alpha,-)} \forall x_1 ... x_k \left(R_\alpha(\overline{x}) \to \exists y \left(R_\beta(y, x_2, ... x_k)\right)\right)$

Clearly, T^G is an $\forall\exists$-theory, and although we have not gone to the trouble, it *is* possible to give a complete set of k-variable $\forall\exists$-axioms for T^G as well.

There is a pair of relational polynomial-time computable transformations

$$-^G : \mathbf{fin}[T] \longrightarrow \mathbf{fin}[T^G], \; -^{\mathbf{mod}} : \mathbf{fin}[T^G] \longrightarrow \mathbf{fin}[T]$$

which completely characterize the relationship between T and T^G.[4] Firstly, suppose $\mathcal{M} \in \mathbf{fin}[T]$; we define $\mathcal{M}^G$ to be the ρ^G-structure with universe M and the obvious interpretations,

$$R_\alpha^{\mathcal{M}^G} = \left\{\overline{a} \in M^k : tp^k(\overline{a};\mathcal{M}) = \alpha\right\}$$

[4] Again, here and in the sequel, $\mathbf{fin}[T]$ and $\mathbf{fin}[T^G]$ denote the classes of finite models of T and T^G, respectively. Also, for a formal definition of relational polynomial-time, see chapter 3.

for each $\alpha \in S^k_k(T)$. The fact that $\mathcal{M}^G \models T^G$ is an easy consequence of theorem 2.1 or of theorem 2.5. Secondly, suppose $\mathfrak{A} \in \mathbf{fin}[T^G]$ with universe A. For $R^{(r)} \in \rho$, we set

$$R^{\mathfrak{A}^{\mathbf{mod}}} = \left\{(a_{i_1}, ..., a_{i_r}) : (a_1, ..., a_k) \in R^{\mathfrak{A}}_{\alpha}, T \models \forall \overline{x}\, (\alpha(\overline{x}) \rightarrow R(x_{i_1}, ...x_{i_r}))\right\}$$

It is essentially trivial to show that the ρ-structure $\mathfrak{A}^{\mathbf{mod}}$ is well-defined and, indeed, a model of T. Collecting these facts, we have:

Observation. *The transformations $-^G$ and $-^{\boldsymbol{mod}}$ are inverses of each other; that is to say, for any $\mathcal{M} \in \boldsymbol{fin}[T]$ and any $\mathfrak{A} \in \boldsymbol{fin}[T^G]$, $(\mathcal{M}^G)^{\boldsymbol{mod}} = \mathcal{M}$ and $(\mathfrak{A}^{\boldsymbol{mod}})^G = \mathfrak{A}$.*

A model $\mathfrak{A}$ of T^G is called a *game tableau for T*, and the model $\mathfrak{A}^{\mathbf{mod}}$ is sometimes called the *realization* of $\mathfrak{A}$. Moreover, for a given model $\mathcal{M}$ of T, the structure $\mathcal{M}^G$ is called the *game tableau of* $\mathcal{M}$; thus, a model $\mathcal{M}$ of T is the unique realization of its own game tableau. The theory T^G is the *theory of game tableaux of* T. Abusing notation slightly, we will write $T^G_\forall$ for the sub-theory consisting of the axioms G1 through G4.

Since there is nothing interesting to distinguish a model of T from its game tableau and the transformation is polynomial-time computable, we generally will not distinguish between finite models of T and finite models T^G; consequently, we will also dispense with the gothic script. Somewhat strangely, we will take as assumed that a model of $\mathcal{M}$ of T is rendered as a model T^G – that is, as its own game tableau. In the next section, we will see that working with game tableaux makes a model-theoretic analysis much more tractable than would be the case in the original signature. The correspondence goes just a bit further in the following proposition (whose proof we omit because it is very simple).

Proposition 2.6. *Let $\mathcal{M}$ and $\mathcal{N}$ be models of T, and let $A \subseteq M$. For any mapping $f : A \rightarrow N$, the following are equivalent:*

1. *f is a partial L^k-elementary embedding $\mathcal{M} \rightharpoonup \mathcal{N}$.*

2. *f is a partial ρ^G-isomorphism $\mathcal{M}^G \rightharpoonup \mathcal{N}^G$.*

In particular, if $\mathcal{M}$ is a model of T, then the complete *quantifier-free* type $qtp(\overline{a}; \mathcal{M}^G)$ of a tuple $\overline{a}$ in the sense of $\mathcal{M}^G$ is equivalent, for our purposes, to the complete k-variable type $tp^k(\overline{a}; \mathcal{M})$.

2.2.1 Capped theories and amalgamation in $\mathbf{fin}[T^G]$

We will say that T is a *capped theory* if for any finite model $\mathcal{A}$ of $T^G_\forall$, there is a finite model $\mathcal{G}$ of T^G such that $\mathcal{A} \leq \mathcal{G}$ – that is, such that $\mathcal{A}$ is an induced substructure of $\mathcal{G}$. For our study, the assumption that T is capped is both natural and rather weak.

Lemma 2.7. *Suppose $\mathcal{M}_0$ and $\mathcal{M}_1$ are models of T^G. Suppose $\mathcal{A}$ is a substructure of both $\mathcal{M}_0$ and $\mathcal{M}_1$, and $M_0 \cap M_1 = A$. (In particular, $\mathcal{A} \models T^G_\forall$.) Then there is a model $\mathcal{C}$ of $T^G_\forall$ and ρ^G-embeddings $g_i : \mathcal{M}_i \rightarrow \mathcal{C}$ such that $g_0{\restriction}A = g_1{\restriction}A$.*

Proof. The idea of the proof is to construct a sort of free-join of $\mathcal{M}_0$ and $\mathcal{M}_1$ over A. It will not be a genuine free-join because even $T^G_\forall$ may induce some additional equalities of elements, and the modified equality relation will then be a non-trivial equivalence relation, say E, on $M_0 \cup M_1$. (We call this the amalgam's *model of equality.*) It's key, then, to maintain the condition $E \cap (M_i \times M_i) = 1_{M_i}$, $i < 2$, in order to avoid obstructing the embeddings. It turns out that maintaining this invariant through the construction is actually sufficient to obtain the amalgam over A.

Let $Z = M_0 \cup M_1$, and let $Q^0_\alpha = R^{\mathcal{M}_0}_\alpha \cup R^{\mathcal{M}_1}_\alpha$ for each $\alpha \in S^k_k(T)$. Furthermore, set $X_0 = Z^k \setminus \bigcup_\alpha Q^0_\alpha$ and $E_0 = 1_{M_0} \cup 1_{M_1}$. Suppose we are then given,

$$X_s \subsetneq X_{s-1} \subsetneq \cdots \subsetneq X_0$$

such that if $\overline{c} \in X_s$ and $\sigma \in \mathit{Sym}[k]$, then $(c_{\sigma(1)}, ..., c_{\sigma(k)}) \in X_s$, and

$$E_s \supsetneq E_{s-1} \supsetneq \cdots \supsetneq E_0$$

where E_s is an equivalence relation on Z such that $E_s \cap (M_i \times M_i) = 1_{M_i}$ for $i = 0, 1$. Let $0 < t < k$, and let $c_1, ..., c_t \in M_0$ and $c_{t+1}, ..., c_k \in M_1$ such that $\overline{c} \in X_s$. Let $\eta_0, \eta_1 \in S^k_k(T)$ such that

$$\mathcal{M}^{\mathbf{mod}}_0 \models \eta_0(c_1, ..., c_t, c_t, ..., c_t)$$

and

$$\mathcal{M}^{\mathbf{mod}}_1 \models \eta_1(c_{t+1}, ..., c_k, c_k, ..., c_k)$$

For brevity, we identify $\eta_0(\overline{x})$, which asserts $\bigwedge^k_{i=t+1} x_i = x_t$, with the t-type it asserts on $x_1, ..., x_t$, and similarly for η_1. We then take the following actions:

1. Let $\alpha \in S^k_k(T)$ such that

   $$T \models \forall \overline{x}\, (\alpha(\overline{x}) \to \eta_0(x_1, ..., x_t) \wedge \eta_1(x_{t+1}, ..., x_k))$$

   Set

   $$Q^{s+1}_\alpha = Q^s_\alpha \cup (\eta_0(M^t_0) \times \eta_1(M^{k-t}_1))$$

   defining $Q^{s+1}_{\alpha^\sigma}$ analogously for each $\sigma \in \mathit{Sym}[k]$.

2. Let E_{s+1} be the $\subseteq$-minimal equivalence relation on Z containing E_s and each (c_i, c_j), $i \leq t < j$, such that $T \models \forall \overline{x}\, (\alpha(\overline{x}) \to x_i = x_j)$.

Claim. *We can choose α so that $E_{s+1} \cap (M_i \times M_i) = 1_{M_i}$, $i = 0, 1$.*

proof of claim. We prove the claim for $i = 0$; the other statement follows by symmetry. Note that we may assume $s > 0$. Suppose $a, b \in M_0$, $a \neq b$ and $aE_{s+1}b$. We may assume that $(a, b) \in E_{s+1} \setminus E_s$ and that aE_sc and $bE_{s+1}c$ for some $c \in M_1$. In particular, there are (w.l.o.g.) elements

$$\begin{aligned} a_1 = a, a_2, ..., a_{t'} &\in M_0 \\ c'_{t'+1} = c, c'_2, ..., c_{k-t'} &\in M_1 \\ b_1 = b, b_2, ..., b_t &\in M_0 \\ c_{t+1} = c, c_2, ..., c_{k-t} &\in M_1 \end{aligned}$$

such that at step $s-1$, we acted on

$$\zeta_0 = tp^k(pad_k(\overline{a}); \mathcal{M}_0^{\mathbf{mod}}),\ \zeta_1 = tp^k(pad_k(\overline{c}'); \mathcal{M}_1^{\mathbf{mod}})$$

and at step s (as above), we acted on

$$\eta_0 = tp^k(pad_k(\overline{b}); \mathcal{M}_0^{\mathbf{mod}}),\ \eta_1 = tp^k(pad_k(\overline{c}); \mathcal{M}_1^{\mathbf{mod}})$$

Since $\zeta_0 \wedge \zeta_1 \models x_1 = x_{t'+1}$ and $\eta_0 \wedge \eta_1 \models x_1 = x_{t+1}$, we now that

$$tp^k(a; \mathcal{M}_0^{\mathbf{mod}}) = tp^k(c; \mathcal{M}_1^{\mathbf{mod}}) = tp^k(b; \mathcal{M}_0^{\mathbf{mod}})$$

As $\mathcal{M}_0$ is a model of T^G, there are $a'_1 = a, a'_2, ...a'_t \in M_0$ such that

$$tp^k(pad_k(\overline{a}'); \mathcal{M}_1^{\mathbf{mod}}) = \eta_0$$

Again, because $\mathcal{M}_0$ is a model of T^G, there are $d_{t+1}, ..., d_k \in M_0$ such that $tp^k(pad_k(\overline{d}); \mathcal{M}_0^{\mathbf{mod}})$ is equal to η_1. Now,

$$\begin{aligned} \eta_0(\overline{a}') \wedge \eta_1(\overline{d}) &\Rightarrow a = d_{t+1} \\ \eta_0(\overline{b}) \wedge \eta_1(\overline{d}) &\Rightarrow b = d_{t+1} \end{aligned}$$

so in fact, $a = b$, a contradiction. □

Since Z is finite, there is a number $n < \omega$ such that $X_n = \emptyset$. (In fact, $n \leq |S_k^k(T)|^2$.) Let $C = Z/E_n$, and for $\alpha \in S_k^k(T)$, let

$$R_\alpha^{\mathcal{C}} = \{(c_1/E_n, ..., c_k/E_n) : (c_1, ..., c_k) \in Q_\alpha^n\}$$

For $i = 0, 1$, define $g_i : B_i \to C$ by $g_i(b) = b/E_n$. It remains to verify that the triple $(\mathcal{C}, g_0, g_1)$ satisfies the requirements of the lemma.

G1: For each k-tuple $\overline{c} = (c_1, ..., c_k) \in Z^k$, either $\overline{c} \in M_0^k \cup M_1^k$ or $\overline{c} \in X_{s-1} \setminus X_s$ for some unique $s \leq n$; hence, $\overline{c}$ is certainly assigned a unique type.

G2: Immediate from the claim we proved above.

G3: Immediate from the construction.

G4: Immediate.

It's then relatively easy to see that g_0 and g_1 are ρ^G-embeddings that agree on A (in fact, each is the identity map on A). □

The lemma, together with the assumption that T is a capped theory, easily yields the following very useful fact.

Theorem 2.8 (AP/sets in $\mathbf{fin}[T^G]$). *Assume that T is capped. Suppose $\mathcal{M}_0$ and $\mathcal{M}_1$ are models of T^G. Suppose $\mathcal{A}$ is a substructure of both $\mathcal{M}_0$ and $\mathcal{M}_1$, and $M_0 \cap M_1 = A$. Then there is a model $\mathcal{N}$ of T^G and ρ^G-embeddings $g_i : \mathcal{M}_i \to \mathcal{N}$ such that $g_0{\restriction}A = g_1{\restriction}A$.*

2.2.2 Example: Vector spaces over finite fields

In this aside, we will see that G-amalgamation and amalgamation in the sense of the original signature do not coincide. In fact G-amalgamation is a strictly weaker assumption than the latter. For our example, fix a finite field $\mathbb{F} = GF(p^n)$ for some prime number p. The signature of interest is, then,

$$\rho = \left\{ R_+^{(3)}, 0, (R_a^{(2)} : a \in \mathbb{F}) \right\}$$

Assume $k \geq 10$, and let $\mathcal{M}_0$ be an $\mathbb{F}$-vector space of dimension no less than k^2. (Later in this chapter, we will see that the dimensions $k-1$ and $d \geq k-1$ cannot be distinguished by the k-variable logic.) We make $\mathcal{M}_0$ into a ρ-structure by interpreting the relation symbols as follows:

$$R_+^{\mathcal{A}_0} = \left\{ (u, v, w) \in M_0^3 : u + v = w \right\}$$

and

$$R_a^{\mathcal{A}_0} = \{ (u, au) : u \in M_0 \}$$

for each $a \in \mathbb{F}$. Let $\mathtt{Vect}_{\mathbb{F}}^k = Th^k(\mathcal{M}_0)$. It's not terribly difficult to verify that since $k \geq 10$ and $\dim(\mathcal{A}_0) \geq k^2$, every model of $\mathtt{Vect}_{\mathbb{F}}^k$ is an $\mathbb{F}$-vector space, and $\mathtt{Vect}_{\mathbb{F}}^k$ has infinitely many finite models up to isomorphism. To make these observations somewhat more precise, we have the following little theorem:

Theorem 2.9. *Let $\mathcal{M}_0$ and $\mathcal{M}_1$ be vector spaces over $\mathbb{F}$ rendered as ρ-structures as above. The following are equivalent:*

1. $\mathcal{M}_0 \equiv^k \mathcal{M}_1$
2. *One of the following holds:*
   - $\dim(\mathcal{M}_0) = \dim(\mathcal{M}_1) < k$;

- $\dim(\mathcal{M}_0) \geq k$ *and* $\dim(\mathcal{M}_1) \geq k$.

To prove the theorem, it's necessary to prove a preliminary technical lemma characterizing winning strategies for the Duplicator, and in the end, we prove a more convenient restatement.

Fix a finite field $\mathbb{F}$, and let $k \geq 10$. Let $m, n < \omega$, and suppose σ is some strategy for the Duplicator in $PG_k(\mathbb{F}^m, \mathbb{F}^n)$, *a priori* not necessarily a winning strategy. For $d \leq k$, we say that σ *is d-linear* if at the end of every round of any game she plays according to σ, the partial ρ-isomorphism $\{f_A(i), f_B(i)) : i \in D\}$ extends uniquely to an isomorphism of $\mathbb{F}$-vector spaces $\langle rng(f_A)\rangle \to \langle rng(f_B)\rangle$ (where $A = \mathbb{F}^m$ and $B = \mathbb{F}^n$). It's easy enough to see that if σ is a $(k-1)$-linear strategy, then σ is a winning strategy. Conversely, we have:

Lemma 2.10. *If σ is a winning strategy for the Duplicator in $PG_k(\mathbb{F}^m, \mathbb{F}^n)$, then σ is $(k-1)$-linear.*

Proof. For a contradiction, suppose σ is not $(k-1)$-linear. In some play of the game, there is a first round, say round r, in which the Duplicator neglects to maintain the $(k-1)$-linearity condition. Without loss of generality, we may assume that the scenario at the end of this round is as follows:

- $f_A(i) = a_i$ for each $i \in [k]$ and $a_k \notin \langle a_2, ..., a_{k-1}\rangle$.
- $f_B(i) = a_i$ for each $i \in [k]$ and $b_k \in \langle b_2, ..., b_{k-1}\rangle$, where $b_k = \sum_{i=2}^{k-1} \beta_i b_i$ and $\beta_2, ..., \beta_{k-1}$ are elements of the field.

We show that from this point, the Spoiler can guarantee a win. The first step is to replace $f_A(i) = a_i$ by $f_A(i) = \hat{a}_i = \beta_i a_i$ for each $i = 2, ..., k-1$ so that the Duplicator is forced to set $f_B(i) = \hat{b}_i = \beta_i b_i$ as well (or lose immediately). For $i = 2, ..., k-1$, the Spoiler enforces the following protocol.

- Round $r + 2(i-2) + 1$. The Spoiler sets $f_A(1) := \hat{a}_i$. Since $\mathbb{F}^m \models R_{\beta_i}(f_A(i), f_A(1))$, the Duplicator must respond by setting $f_B(1) := \beta_i \cdot f_B(i) = \hat{b}_i$.
- Round $r+2(i-2)+2$. The Spoiler sets $f_A(i) := \hat{a}_i$. Again, since $\mathbb{F}^m \models (f_A(1) = f_A(i))$, the Duplicator must now set $f_B(i) = \hat{b}_i$.

At the end of round $r + 2(k-3) + 2$, the situation is as follows:

1. $f_A(1) = \hat{a}_2$, $f_A(i) = \hat{a}_i$ for $i = 2, ..., k-1$ and $f_A(k) = a_k$
2. $f_B(1) = \hat{b}_2$, $f_B(i) = \hat{b}_i$ for $i = 2, ..., k-1$ and $f_B(k) = b_k$

Next, the Spoiler sets $f_A(1) := f_A(2) = \hat{a}_2$, so that Duplicator must respond with $f_B(1) := f_B(2) = \hat{b}_2$. Now, with $r' = r + 2(k-3) + 3$, the Spoiler enforces the following protocol for $i = 3, ..., k-1$:

- Round $r' + 2(i-3) + 1$. The Spoiler resets $f_A(2) := f_A(1)$, so the Duplicator must set $f_B(2) := f_B(1)$.
- Round $r' + 2(i-3) + 2$. The Spoiler resets $f_A(1) := f_A(2) + f_A(i) = \sum_{j=2}^{i} \hat{a}_i$, so again, the duplicator must set $f_B(1) := f_B(2) + f_B(i) = \sum_{j=2}^{i} \hat{b}_j$

Thus, at end of round $r'' = r' + k - 4$, we have

$$f_A(1) = \sum_{i=2}^{k-1} \hat{a}_i \neq a_k = f_A(k) \text{ and } f_B(1) = \sum_{i=2}^{k-1} \hat{b}_i = b_k = f_B(k)$$

and the Duplicator has lost, contradicting the assumption that σ was a winning strategy. □

We can now prove a slightly more convenient (though obviously equivalent) restatement of theorem 2.9.

Proposition 2.11. *Let $\mathbb{F}$ be a finite field, and let $k \geq 10$. Let $m, n < \omega$.*

1. *If $m, n < k$, then $\mathbb{F}^m \equiv^k \mathbb{F}^n$ if and only if $m = n$.*
2. *If $m < k \leq n$, then $\mathbb{F}^m \not\equiv^k \mathbb{F}^n$.*
3. *If $m, n \geq k$, then $\mathbb{F}^m \equiv^k \mathbb{F}_n$.*

Proof of 1 and 2. If $m = n$, then $\mathbb{F}^m = \mathbb{F}^n$, so there is nothing to do. If $m < k-1$ and $m < n$, then there is obviously no hope of finding a $(k-1)$-linear strategy in $PG_k(\mathbb{F}^m, \mathbb{F}^n)$. □

Proof of 3. Since $\equiv^k$ is evidently transitive, it suffices to show that for any $m > k$, then the Duplicator has a $(k-1)$-linear strategy in $PG_k(\mathbb{F}^m, \mathbb{F}^k)$. To prove this, it suffices to deal with the following scenario:

1. $\{f_B(1), ..., f_B(k)\}$ is a basis of $\mathbb{F}^k = B$, and $\{f_A(1), ..., f_A(k)\}$ is linearly independent in $\mathbb{F}^k = A$.
2. $\{(f_A(i), f_B(i)) : i \in [k]\}$ extends uniquely to an isomorphism $\langle rng(f_A)\rangle \to \mathbb{F}^k$ of vector spaces.
3. The Spoiler chooses $f_A(k) \notin \langle f_A(1), ..., f_A(k-1)\rangle$.

Now, the Duplicator can choose any

$$f_B(k) \in \mathbb{F}^k \setminus \bigcup \left\{ \langle f_B(i)\rangle_{i \in X} : X \in \binom{[k-1]}{k-2} \right\}$$

and thereby maintain the $(k-1)$-linearity condition. □

Observation. $\mathtt{Vect}^k_{\mathbb{F}}$ *does not have amalgamation over sets with respect to* ρ.

Proof. (This example was worked out in [3].) Let $\mathcal{M}_0$ and $\mathcal{M}_1$ be $\mathbb{F}$-vector spaces of dimension $\geq k^2$, understood as ρ-structures, and let A be a set of exactly k elements such that $M_0 \cap M_1 = A$. Assume that $\sum A = 0$ in $\mathcal{M}_0$, and $\sum A = b \neq 0$ in $\mathcal{M}_1$. Further assume that for any $X \in \binom{A}{k-2}$, $\sum X \neq 0$ in both $\mathcal{M}_0$ and $\mathcal{M}_1$. One can show that under this condition $tp^k(A; \mathcal{M}_0)$ and $tp^k(A; \mathcal{M}_1)$ are identical. Towards a contradiction, suppose $\mathcal{N}$ is an $\mathbb{F}$-vector space and $g_i : \mathcal{M}_i \to \mathcal{N}$ are L^k-elementary embeddings such that $g_0 \restriction A = g_1 \restriction A$. Clearly, g_0 and g_1 are embeddings of vector spaces ($\mathbb{F}$-linear monomorphisms) in the usual sense. Let $X \in \binom{A}{k-2}$, let $c_1 = \sum X$, and $A \setminus X = \{c_2, c_3\}$. Since $c_1 + c_2 + c_3 = 0$ in $\mathcal{M}_0$ and $k \geq 10$, we know that

$$g_0(c_1) + g_0(c_2) + g_0(c_3) = g_0(0) = 0$$

On the other hand,

$$g_1(c_1) + g_1(c_2) + g_1(c_3) = g_1(b) \neq 0$$

It can also be shown that $g_0 \restriction A = g_1 \restriction A$, together with the fact that $|X| \leq k-2$, implies that $g_0(c_1) = g_1(c_1)$, and by definition $g_0(c_i) = g_1(c_i)$, $i = 2, 3$, just because $c_2, c_3 \in A$. Therefore,

$$0 = g_0(c_1 + c_2 + c_3) = g_1(c_1 + c_2 + c_3) \neq 0$$

which is impossible. $\square$

Observation. $\mathtt{Vect}^k_{\mathbb{F}}$ *is capped. In particular,* $\mathtt{Vect}^k_{\mathbb{F}}$ *has* G*-amalgamation.*

Proof. Suppose $\mathcal{A}$ is a model of $(\mathtt{Vect}^k_{\mathbb{F}})^G_{\forall}$. There is a subset $I \subseteq S^k_k(\mathtt{Vect}^k_{\mathbb{F}})$ comprising those types $\alpha(x_1, ..., x_k)$ which *do not* assert that $x_{i_1}, ..., x_{i_{k-2}}$ are linearly dependent for some $i_1 < \cdots < i_{k-2} \leq k$. Let $B \subseteq A$ be a $\subseteq$-maximal subset of A subject to the requirement that for all $b_1, ..., b_k \in B$, if the b_i's are pairwise distinct and $(b_1, ..., b_k) \in R^{\mathcal{A}}_{\alpha}$, then $\alpha \in I$. Observe that if $c \in A \setminus B$, then for some $\{b_1, ..., b_{k-1}\} \in \binom{B}{k-1}$, there is an $\alpha \notin I$ such that $\bar{b}c \in R^{\mathcal{A}}_{\alpha}$; for otherwise, B would not be $\subseteq$-maximal. Let $d = |B|$, and let $\mathcal{M}$ be an $\mathbb{F}$-vector space of dimension $\geq d$. Let X be a basis of $\mathcal{M}$. It is not difficult to see that there is an injection $g_0 : B \to X$ which is simultaneously a ρ^G-embedding $\mathcal{A} \restriction B \to \mathcal{M}^G$. It is similarly easy to verify that g_0 extends uniquely to a ρ^G-embedding $g : \mathcal{A} \to \mathcal{M}^G$, and this completes the demonstration. $\square$

We have seen, then, a natural example of capped L^k-theory which nonetheless does not admit amalgamation over sets with respect to its original signature. Thus, the passage from T to T^G is nontrivial.

The final theorem of this section turns out to be unavoidable (though terribly convenient) in a few arguments, although its use in those cases does leave us with a feeling of bad conscience.

Theorem 2.12. *If T is capped, then $\boldsymbol{fin}[T^G]$ has a Fraššé limit. That is, there is a model $\mathfrak{M}$ of T^G which satisfies:*

1. *$Th(\mathfrak{M})$ is $\aleph_0$-categorical and eliminates quantifiers.*
2. *$\mathfrak{M}$ is ultrahomogeneous and $Age(\mathfrak{M}) = \boldsymbol{fin}[T^G_\forall]$*
3. *If $\mathcal{M} \in \boldsymbol{fin}[T^G]$, $A \subseteq M$, and $g_0 : (A; \mathcal{M}) \to \mathfrak{M}$ is a partial ρ^G-isomorphism, then there is a ρ^G-isomorphism $g : \mathcal{M} \to \mathfrak{M}$ extending g_0.*

Proof. It is straightforward to verify that when T is capped, $\mathbf{fin}[T^G_\forall]$ is hereditary (has HP) and has the joint-embedding and amalgamation properties (JEP and AP, respectively). By material of chapter 7 of [15], we know that $\mathbf{fin}[T^G_\forall]$ has a Fraïssé limit $\mathfrak{M}$. Using again the fact that T is capped, one can easily show that there is an ascending chain of finite substructures of $\mathfrak{M}$, say

$$\mathcal{M}_0 < \cdots < \mathcal{M}_n < \cdots <_{\mathbf{fin}} \mathfrak{M}$$

such that $\mathcal{M}_n \vDash T^G$ for each $n < \omega$. Since T^G is $\forall\exists$-theory, it follows that $\mathfrak{M} = \bigcup_{n<\omega} \mathcal{M}_n$ is also a model of T^G, and this completes the proof of the theorem. □

2.3 Closure, coherent sequences and indiscernibles

For purely stylistic reasons, we will usually (with a few exceptions) eschew using infinite limit models in many of our arguments, but we will still need to use some of the machinery of indiscernible sequences, which should come as little surprise to the reader with a background in model theory. In this section, we introduce an adaptation of this technology developed in [17] for use when working with classes of finite models – even classes without any sort of amalgamation property and under fairly general notions of "logic." In all of our analyses, we will use the "logic" consisting of quantifier-free formulas with respect to a finite relational signature σ (e.g. ρ^G as derived from a fixed finite-variable theory); the key properties of this logic pertaining to indiscernible sequences is the following. Let K be a class of finite ρ-structures which is closed under isomorphism. We first note some useful properties of the logic of quantifier-free formulas:

($\star$) There is a function $t_\sigma : \omega \times \omega \to \omega$ such that

$$|S^{\mathrm{qf}}_n(A; \mathcal{M})| \leq t_\sigma(|A|, n)$$

for all $\mathcal{M} \in K$, $A \subseteq M$ and $0 < n < \omega$.

($\star\star$) There is a number n_ρ such that for all $n \geq n_\rho$, $\mathcal{M} \in K$, $A \subseteq M$ and $b_1, ..., b_n \in M$,

$$qtp(\overline{b}/A; \mathcal{M}) = \bigcup_{i_1,...,i_{n_\sigma} \in [n]} qtp(b_{i_1}, ..., b_{i_{n_\sigma}}/A; \mathcal{M})$$

We now move to the essential definitions. Let $\mathcal{M} \in K$, $C \subseteq M$, and $n < \omega$. An *(n-ary) coherent sequence over* $(C; \mathcal{M})$ is a sequence of pairs $(\mathcal{A}_i, \overline{a}_i)_{i<\omega}$ such that for all $i < \omega$,

- $\mathcal{A}_i \in K_C$ and $\overline{a}_0, ..., \overline{a}_i \in A_i^n$
- If $j \leq i$, then $qtp(\overline{a}_0...\overline{a}_j/C; \mathcal{A}_j) = qtp(\overline{a}_0...\overline{a}_j/C; \mathcal{A}_i)$

The first lemma simply asserts the existence of coherent sequences where they are needed; the proof is a relatively routine application of König's lemma using property ($\star$).

Lemma 2.13. *Let $\mathcal{M} \in K$, $C \subseteq M$, and $n < \omega$. Let $\Gamma = (\mathcal{B}_m, (\overline{b}_i^{(m)})_{i \leq m})_{m<\omega}$ be a sequence of pairs such that $\mathcal{B}_m \in K_C$ and $\overline{b}_i^{(m)} \in B_m^n$ for all $i \leq m < \omega$. Then, there is a coherent sequence $(\mathcal{A}_j, \overline{a}_j)_{j<\omega}$ over $(C; \mathcal{M})$* strictly patterned *on Γ in the following sense: For each $j < \omega$, there is an $m_j < \omega$ such that $qtp(\overline{a}_0...\overline{a}_j/C; \mathcal{A}_j) = qtp(\overline{b}_0^{(m_j)}...\overline{b}_j^{(m_j)}/C; \mathcal{B}_{m_j})$. where $qtp(\overline{a}/C; \mathcal{M})$ denotes the complete quantifier-free type over $\overline{a}$ over C in the sense of $\mathcal{M}$.*

Now, fix an n-ary coherent sequence $(\mathcal{A}_i, \overline{a}_i)_{i<\omega}$ over $(C; \mathcal{M})$ for some $\mathcal{M} \in K$ and $C \subseteq M$, and let $C_0 \subseteq C$. We say that $(\mathcal{A}_i, \overline{a}_i)_{i<\omega}$ is *(order) indiscernible over* $(C_0; \mathcal{M})$ (or more succinctly, that it is C_0-indiscernible) if

$$qtp(\overline{a}_{i_1}...\overline{a}_{i_n}/C_0; \mathcal{A}_s) = qtp(\overline{a}_{j_1}...\overline{a}_{j_n}/C_0; \mathcal{A}_t)$$

whenever $n, s, t < \omega$, $i_1, ..., i_n \leq s$, $j_1, ..., j_n \leq t$, and for all $p, q \in [n]$, $i_p < i_q$ if and only if $j_p < j_q$. The next lemma is proved by a straightforward application of Ramsey's theorem together with property ($\star\star$).

Lemma 2.14. *Let $\mathcal{M} \in K$, $C \subseteq M$, and $0 < n < \omega$, and let $(\mathcal{A}_i, \overline{a}_i)_{i<\omega}$ be an n-ary coherent sequence over $(C; \mathcal{M})$. Then, there is an infinite subset $X \subseteq \omega$ such that $(\mathcal{A}_i, \overline{a}_i)_{i \in X}$ is order indiscernible over $(C; \mathcal{M})$.*

Combining these lemmas, we obtain the main proposition of this section, which will used repeatedly throughout this dissertation.

Proposition 2.15. *Let $\mathcal{M} \in K$ and $C \subseteq M$, and let $0 < n < \omega$. Suppose that $\Gamma = (\mathcal{B}_m, (\overline{b}_i^{(m)})_{i \leq m})_{m<\omega}$ is a sequence of pairs such that $\mathcal{B}_m \in K_C$ and $\overline{b}_i^{(m)} \in B_m^n$ for all $i \leq m < \omega$. Then there is a C-indiscernible coherent sequence $(\mathcal{A}_j, \overline{a}_j)$ over $(C; \mathcal{M})$ strictly patterned on Γ as above.*

There is one last bit of notation which we must introduce before proceeding. Suppose $\mathcal{M} \in K$, $A, C \subseteq M$, and $p(x) = qtp(b/C; \mathcal{M})$. Recall that $p(x)$ is K-*algebraic* just in case there is a number $r < \omega$ such that for every $\mathcal{N} \in K_C$, we have $|\{b' \in N : \mathcal{N} \vDash p(b')\}| \leq r$. We also define

$$\kappa(C; \mathcal{M}) = \{a \in M : qtp(a/C; \mathcal{M}) \text{ is } K\text{-algebraic}\}$$

Since we do not have access to quantifiers, it is not *a priori* obvious that κ captures the usual model-theoretic notion of the algebraic-closure of a set. Firstly, it is certainly plausible that K contains a pair of models $\mathcal{M} < \mathcal{M}'$ with $C \subseteq M$ such that $\kappa(C; \mathcal{M}) \subsetneqq \kappa(C; \mathcal{M}')$. Secondly, it is not immediately clear that κ is a closure operator, even when restricted to subsets of a single model. To combat that latter issue, we simply iterate the construction to obtain a closure operator:

$$\begin{aligned} \mathrm{cl}_0(A; \mathcal{M}) &= A \\ \mathrm{cl}_{n+1}(A; \mathcal{M}) &= \mathrm{cl}_n(A) \cup \kappa(\mathrm{cl}_n(A; \mathcal{M}); \mathcal{M}) \\ \mathrm{cl}(A; \mathcal{M}) &= \bigcup_{n<\omega} \mathrm{cl}_n(A; \mathcal{M}) = \bigcup_{n \leq |M|} \mathrm{cl}_n(A; \mathcal{M}) \end{aligned}$$

Now, the following is easily verified:

Observation. *1. If $\mathcal{M} \leq \mathcal{M}'$ in K and $A \subseteq M$, then $\kappa(A; \mathcal{M}) \subseteq \kappa(A; \mathcal{M}')$ and $cl(A; \mathcal{M}) \subseteq cl(A; \mathcal{M}')$.*

2. *There is a function $f_{cl} : \omega \to \omega$ such that if $(\mathcal{M}_i)_{i<\alpha}$, $\alpha \leq \omega$, is an ascending chain in K, $A \subseteq M_0$ and $cl(A; \mathcal{M}_i) \subsetneqq cl(A; \mathcal{M}_{i+1})$ whenever $i + 1 < \alpha$, then $\alpha \leq f_{cl}(|A|)$.*

3. *Suppose $\mathfrak{M}$ is the Fraïssé limit of K, and let $\mathcal{M} \in Age(\mathfrak{M})$ and $A \subseteq M$. Then*

$$cl(A; \mathcal{M}) \subseteq acl^{\mathfrak{M}}(A) = \bigcup \{cl(A; \mathcal{N}) : \mathcal{N} \in Age(\mathfrak{M}), A \subseteq N\}$$

and $acl^{\mathfrak{M}}(A) \cap M = cl(A; \mathcal{M})$.

As an example, we characterize (without proof) the behavior of the cl-operator for vector spaces over finite fields. It is a straightforward consequence of the characterization of winning strategies as $(k-1)$-linear strategies.

Observation. *Let $\mathbb{F}$ be a finite field, and let T be the k-variable theory of finite $\mathbb{F}$-vector spaces of "large" (i.e. $\geq k$) dimension. Let $\mathcal{M} \models T$ and $A_0 \subseteq A \subseteq M$. Then*

$$cl(A; \mathcal{M}^G) = \bigcup \left\{ cl(B; \mathcal{M}^G) = \langle B \rangle : B \in \binom{A}{k-1} \right\}$$

Furthermore, if $\mathfrak{A} = (A; \mathcal{M}^G) \leq \mathcal{M}^G$, then

$$cl(A_0; \mathfrak{A}) = \bigcup \left\{ cl(A_0; \mathfrak{A}) = A_0 \cap \langle B \rangle : B \in \binom{A_0}{k-1} \right\}$$

Chapter 3

Essentials of transformations of structures

In this chapter, we introduce the basic algorithmic machinery of our analysis, deferring more complex characterizations of efficient programs to the following chapters. The ideas on programming languages here are taken almost directly from [8] and [1]. The material on fixed-point logics is relatively standard in finite-model and can be found in [19, 11], first appearing (to our knowledge) in [1]. However, its development to accommodate object-creating operations like our `invent`-operator seems to be new, and the idea of the model-building problem and coherent solutions are original .

3.1 Fundamental objects

3.1.1 Hereditary lists

The notion of hereditary lists is inherited, with some modification, from [8]. For an arbitrary non-empty set X disjoint from $\mathbb{N}$, we define the set of *hereditarily finite lists* as follows:

$$\begin{aligned} \mathrm{HL}_0[X] &= X \cup \mathbb{N} \\ \mathrm{HL}_{n+1}[X] &= \mathrm{HL}_n[X] \cup (\mathrm{HL}_n[X])^{<\omega} \\ \mathrm{HL}[X] &= \bigcup_{n<\omega} \mathrm{HL}_n[X] \end{aligned}$$

For a positive integer t, we define $\mathrm{HL}^{(t)}[X]$ in the same manner except for the modification:

$$\mathrm{HL}_0^{(t)}[X] = X \cup \{1, ..., t\}$$

In our analyses, it will be convenient to fix a countably infinite set $\mathbf{U}_0$ disjoint from $\mathbb{N}$ and such that $\mathbf{U}_0 \cap \mathrm{HL}[\mathbf{U}_0] = \mathbf{U}_0$. Every structure $\mathcal{A}$, then, will be assumed to satisfy

$A \subseteq \mathrm{HL} = \mathrm{HL}[\mathbf{U}_0]$. At times, it will be convenient to endow HL with some structure of its own; namely, for each $0 < l < \omega$, let $J_l \subseteq \mathrm{HL}^{l+1}$ be the relation consisting of $(l+1)$-tuples of the form

$$(a_1, ..., a_l, \langle a_1, ..., a_l \rangle)$$

For $\mathrm{HL}^{(t)}$, we use only $J_1, ..., J_t$, and we name $1, ..., t$ as constant symbols. If ρ is some finite relational signature and $\mathcal{A}$ is a ρ-structure such that $A \subseteq \mathrm{HL}$, then we obtain an expansion $(\mathrm{HL}, \mathcal{A})$ of HL with signature $\rho \cup \{D^{(1)}\} \cup \{J_l\}_l$, with $D^{(\mathrm{HL},\mathcal{A})} = A$ and $R^{(\mathrm{HL},\mathcal{A})} = R^{\mathcal{A}}$ for all $R \in \rho$. We say that $\mathcal{A}$ is *initial* if $A \subseteq \mathbf{U}_0$, and we define $\mathbf{fin}_{init}[\rho]$ to be the set of all finite initial ρ-structures.

Note that permutations σ of $\mathbf{U}_0$ lift naturally to permutations of HL recursively via

$$\sigma\langle a_1, ..., a_n \rangle = \langle \sigma a_1, ..., \sigma a_n \rangle$$

(where we take $\sigma\langle\rangle = \langle\rangle$); we sometimes write $\hat{\sigma}$ for the permutation of HL induced by $\sigma \in \mathit{Sym}(\mathbf{U}_0)$. If $\mathcal{A}$ is a ρ-structure with $A \subseteq \mathrm{HL}$, then $\mathit{Sym}(\mathbf{U}_0)$ also acts on $\mathcal{A}$ as follows; for $\sigma \in \mathit{Sym}(\mathbf{U}_0)$, $\mathcal{A}^\sigma$ is the structure with universe $\sigma[A]$ and

$$R^{\mathcal{A}^\sigma} = \left\{(\sigma a_1, ..., \sigma a_r) : (a_1, ..., a_r) \in R^{\mathcal{A}}\right\}$$

for each $R \in \rho$.

Now, suppose ρ_{in} and ρ_{out} are both finite relational signatures. An *HL-transformation of type* $\rho_{in} \to \rho_{out}$ is a partial function

$$Q : \mathbf{fin}_{init}[\rho_{in}] \rightharpoonup \mathbf{fin}[\rho_{out}]$$

satisfying the following conditions:

1. The graph of Q is recursively enumerable (up to a recursive enumeration of $\mathbf{U}_0$);
2. If $\mathcal{A} \in \mathit{dom}(Q)$, then $A \subseteq ||Q(\mathcal{A})|| \subseteq \mathrm{HL}[A]$;
3. Q is $\mathit{Sym}(\mathbf{U}_0)$-invariant in the following sense: If $\mathcal{A} \in \mathit{dom}(Q)$, $Q(\mathcal{A}) = \mathcal{B}$ and $\sigma \in \mathit{Sym}(\mathbf{U}_0)$, then $Q(\mathcal{A}^\sigma) = \mathcal{B}^{\hat{\sigma}}$

The following theorem is from [8].

Theorem 3.1. *Suppose Q is an HL-transformation of type $\rho_{in} \to \rho_{out}$. Then there is a partial recursive family ψ of mappings of the form*

$$\psi_{\mathcal{A}} : \mathit{Aut}(\mathcal{A}) \hookrightarrow \mathit{Aut}(Q(\mathcal{A}))$$

for $\mathcal{A} \in \mathit{dom}(Q)$ such that for any $g \in \mathit{Aut}(\mathcal{A})$, $\psi_{\mathcal{A}}(g)(b) = \hat{g}(b)$ for every $b \in ||Q(\mathcal{A})||$. Here, "partial recursive" means that the relation

$$\{(\mathcal{A}, \sigma, \sigma_1) : \psi_{\mathcal{A}}(\sigma)\!\downarrow = \sigma_1\}$$

is recursively enumerable up to some encoding of structures and functions.

3.1.2 Weakly constructible k-variable theories

Once again, we fix a finite relational signature ρ, and we fix a complete L^k-theory T (with infinitely many finite models up to isomorphism) in this signature. T^G is the theory of game tableaux for models of T, and $T^G_\forall$ is the theory consisting of the universal sentences of T^G. K denotes the class of finite models of T^G – i.e. $K = \mathbf{fin}[T^G]$ – and $K_\forall$ the class of finite induced substructures of models of T^G. We also assume that T is capped – that is, for every $\mathcal{A} \in K_\forall$, there is a model $\mathcal{M} \in K$ which has $\mathcal{A}$ as an induced substructure. A *weakly coherent solution of the model-building problem for* T is an HL-transformation of type $\rho^G \to \rho_1$, where $\rho^G \subseteq \rho_1$ which satisfies the following requirements:

1. $K_\forall \subseteq dom(Q)$, and for each $\mathcal{A} \in K_\forall$, $Q(\mathcal{A})\restriction\rho^G$ is a model of T^G;

   Moreover, for any k-type $p(\overline{x})$ over A (with respect to $Q(\mathcal{A})$) and any $\mathcal{M} \in K$, if $Q(\mathcal{A})\restriction\rho^G \leq \mathcal{M}$ and $p(\overline{x})$ is algebraic, then $p(M^k) \subseteq ||Q(\mathcal{A})||^k$.

2. For each $\mathcal{A} \in K_\forall$ and $R \in \rho^G$, $R^{Q(\mathcal{A})} \subseteq ||Q(\mathcal{A})\restriction\rho^G||^k$.

   (In this formulation, we have not formally defined the reduct $\mathcal{B}\restriction\rho^G$ for a ρ_1-structure $\mathcal{B}$; for this formality, $\mathcal{B}\restriction\rho^G$ is the ρ^G-structure with universe $B_0 = \bigcup_{R\in\rho^G} \mathit{fld}(R^{\mathcal{B}})$ and interpretations $R^{\mathcal{B}\restriction\rho^G} = R^{\mathcal{B}}$ for each $R \in \rho^G$.)

3. If $\mathcal{A}, \mathcal{B} \in K_\forall$ and $\mathcal{A} \leq \mathcal{B}$, then $\mathcal{A} \leq Q(\mathcal{A})\restriction\rho^G \leq Q(\mathcal{B})\restriction\rho^G$.

Of course, we say that T is *weakly constructible* just in case T admits a weakly coherent solution of the model-building problem. Note that weak coherence is a purely abstract condition. Once we have introduced the machinery of a programming language, we will make a slightly more specialized definition of coherence corresponding to efficient solutions of the model-building problem.

3.2 The `invent`-programming language

In subsection 3.2.1, we introduce the most primitive instructions of our programming language, called *basic expressions*, and their semantics. For this, we choose to understand first-formulas as operators on the set of structures, and to define the *response* of applying a basic expression to a structure, we define a secondary operator `resp`. We will then, in subsection 3.2.2, extend the programming syntax to included some necessary control structures and extend the definition of the `resp`-operator to accommodate these.

3.2.1 Basic expressions

As usual, fix a finite relational signature ρ_1. For evaluating basic expressions, we must understand all first-order formulas in the language of ρ_1 to be relativized to the predicate D, distinguishing the universe of a finite structure $\mathcal{A}$ such that $A \subseteq \mathrm{HL}$; such a ρ_1-structure is called *situated*. Recall, then, that the D-relativization of a first-order ρ_1-formula is defined as follows:

- $R(x_1, ..., x_n)^D = R(x_1, ..., x_n)$
- $(\varphi \wedge \psi)^D = \varphi^D \wedge \psi^D$
- $(\neg\varphi(x_1, ..., x_n))^D = \bigwedge_{i=1}^n D(x_i) \wedge \neg\, \varphi^D$
- $(\exists x\, \varphi)^D = \exists x(D(x) \wedge \varphi^D)$

(Note that D is never going to be a symbol of ρ_1, neither will be any of the predicate symbols structuring HL itself.) Throughout our analyses, we will suppress the superscript D in our notation, but it is important to keep in mind that all sets in question – the definable sets in particular – are subsets of some cartesian power of D. Now, suppose that $\varphi(x_1, ..., x_n)$ is a first-order formula in the language of ρ_1; the formula φ, then, supports two kinds of operators on situated ρ_1-structures:

$$\begin{aligned} \{\overline{x}|\varphi\}(\mathcal{A}) &= \{(a_1, ..., a_n) \in A^n : \mathcal{A} \vDash \varphi(a_1, ..., a_n)\} \\ \mathtt{invent}_k\{\overline{x}|\varphi\}(\mathcal{A}) &= \{(a_1, ..., a_n, 1\widehat{\ }\overline{a}, ..., (k-n)\widehat{\ }\overline{a}) : (a_1, ..., a_n) \in \{\overline{x}|\varphi\}(\mathcal{A})\} \end{aligned}$$

(In the latter case, we assume that $k > n$.) These operators give rise to the interpretations of the the *basic expressions* of our programming language:

- If $\varphi(x_1, ..., x_n) \in L(\rho_1)$ and $X \in \rho_1$ with $ari(X) = n$, then

  $$(X \leftarrow \{\overline{x}|\varphi\})$$

  is a basic expression. If $\mathcal{A}$ is a situated ρ_1-structure, then

  $$\mathtt{resp}((X \leftarrow \{\overline{x}|\varphi\}), \mathcal{A}) = \mathcal{A}\,[\,X/\{\overline{x}|\varphi\}(\mathcal{A})\,]$$

  That is, $\mathtt{resp}((X \leftarrow \{\overline{x}|\varphi\}), \mathcal{A})$ is the ρ_1-structure with universe A,

  $$Y^{\mathtt{resp}((X \leftarrow \{\overline{x}|\varphi\}), \mathcal{A})} = Y^{\mathcal{A}}$$

  whenever $Y \neq X$, and

  $$X^{\mathtt{resp}((X \leftarrow \{\overline{x}|\varphi\}), \mathcal{A})} = \{\overline{x}|\varphi\}(\mathcal{A})$$

- If $\varphi(x_1, ..., x_n) \in L(\rho_1)$ and $X \in \rho_1$ with $ari(X) = k > n$, then

  $$(X \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\})$$

  is a basic expression. If $\mathcal{A}$ is a situated ρ_1-structure, then

  $$\mathtt{resp}((X \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\}), \mathcal{A}) = \mathcal{A}'$$

  where $A' = A \cup \mathit{fld}(\mathtt{invent}_k\{\overline{x}|\varphi\}(\mathcal{A}))$, $Y^{\mathcal{A}'} = Y^{\mathcal{A}}$ if $Y \in \rho_1 \setminus \{X\}$ and $X^{\mathcal{A}'} = \mathtt{invent}_k\{\overline{x}|\varphi\}(\mathcal{A})$.

3.2.2 Control structures

The basic expressions encode all of the direct actions taken on a structure $\mathcal{A}$, but of course, the programming language also requires control structures. We define the set of `invent`-programs (with respect to ρ_1) by the following grammar:

$$\alpha ::= \text{<basic expressions>} \mid \alpha;\alpha \mid (\mathbf{if}\,\psi\,\mathbf{then}\,\alpha\,\mathbf{else}\,\alpha) \mid (\mathbf{while}\,\psi\,\mathbf{do}\,\alpha)$$

where ψ ranges over sentences of $L(\rho_1)$, relativized to D as always. Less succinctly, `invent`-programs are the members of the following inductively-defined set $\mathbb{P}$ of formal expressions.

1. Every basic expression is in $\mathbb{P}$.
2. If P_1 and P_2 are in $\mathbb{P}$, then $P_1; P_2$ is in $\mathbb{P}$.
3. If P_1 and P_2 are in $\mathbb{P}$ and ψ is D-relativized sentence in $L(\rho_1)$, then $(\mathbf{if}\,\psi\,\mathbf{then}\,P_1\,\mathbf{else}\,P_2)$ is in $\mathbb{P}$.
4. If P is in $\mathbb{P}$ and ψ is D-relativized sentence in $L(\rho_1)$, then $(\mathbf{while}\,\psi\,\mathbf{do}\,P)$ is in $\mathbb{P}$.

The semantics of these programs is defined inductively as follows (our measure of time-complexity is defined in parallel):

- $\texttt{resp}(P_1; P_2, \mathcal{A}) = \texttt{resp}(P_2, \texttt{resp}(P_1, \mathcal{A}))$
  $\texttt{cpx}(P_1; P_2, \mathcal{A}) = \texttt{cpx}(P_1, \mathcal{A}) + \texttt{cpx}(P_2, \texttt{resp}(P_1, \mathcal{A}))$
- $\texttt{resp}((\mathbf{if}\,\psi\,\mathbf{then}\,P_1\,\mathbf{else}\,P_2), \mathcal{A}) = \begin{cases} \texttt{resp}(P_1, \mathcal{A}) & \text{if } \mathcal{A} \models \psi \\ \texttt{resp}(P_2, \mathcal{A}) & \text{if } \mathcal{A} \not\models \psi \end{cases}$
  $\texttt{cpx}((\mathbf{if}\,\psi\,\mathbf{then}\,P_1\,\mathbf{else}\,P_2), \mathcal{A}) = \begin{cases} \texttt{cpx}(P_1, \mathcal{A}) & \text{if } \mathcal{A} \models \psi \\ \texttt{cpx}(P_2, \mathcal{A}) & \text{if } \mathcal{A} \not\models \psi \end{cases}$
- $\texttt{resp}((\mathbf{while}\,\psi\,\mathbf{do}\,P), \mathcal{A}) = \mathcal{A}_{n^*}$ where

  $$\begin{aligned} \mathcal{A}_0 &= \mathcal{A} \\ \mathcal{A}_{i+1} &= \texttt{resp}(P, \mathcal{A}_i) \\ n^* &= \min\{i < \omega : \mathcal{A}_i \models \neg\psi\} \end{aligned}$$

  Of course, such an n^* may not exist or some $\texttt{resp}(P, \mathcal{A}_i)$ may be undefined, and in either case, $\texttt{resp}((\mathbf{while}\,\psi\,\mathbf{do}\,P), \mathcal{A})$ is undefined. Assuming $\texttt{resp}((\mathbf{while}\,\psi\,\mathbf{do}\,P), \mathcal{A})$ is defined, we set

  $$\texttt{cpx}((\mathbf{while}\,\psi\,\mathbf{do}\,P), \mathcal{A}) = \sum_{i<n^*} \texttt{cpx}(P, \mathcal{A}_i)$$

Lemma 3.2. *Let $\mathcal{A}$ be a siturated ρ_1-structure, and let ε be a basic expression over ρ_1. Then, if $0 < r < \omega$ and $\overline{b} \in ||\mathtt{resp}(\varepsilon, \mathcal{A})||^r$, then*

$$\overline{b}^{Aut(\mathcal{A})} = \left\{\hat{\sigma}\overline{b} : \sigma \in Aut(\mathcal{A})\right\} \subseteq ||\mathtt{resp}(\varepsilon, \mathcal{A})||^r$$

Moreover, $\overline{b}^{Aut(\mathcal{A})}$ is a union of orbits of $Aut(\mathtt{resp}(P, \mathcal{A}))$ on r-tuples.

Proof. For the first claim, we give the demonstration when $r = 1$; for $r > 1$, nothing more complicated is going on. If ε is not an `invent`-expression, then obviously, there is nothing to prove. Hence, we assume that ε is of the form

$$(X \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\})$$

where $\varphi = \varphi(x_1, ..., x_n)$ is in $FO[\rho_1]$ and $n < k$. We may assume, then, that there is an n-tuple $(a_1, ..., a_n) \in A^n$ such that $\mathcal{A} \vDash \varphi(a_1, ..., a_n)$ and an $i \in \{1, ..., k - n\}$ such that $b = \langle i, a_1, ..., a_n\rangle$. If $\sigma \in Aut(\mathcal{A})$, then

$$\hat{\sigma}b = \hat{\sigma}\langle i, a_1, ..., a_n\rangle = \langle \sigma i, \sigma a_1, ..., \sigma a_n\rangle = \langle i, \sigma a_1, ..., \sigma a_n\rangle$$

Since σ is an automorphism, we have $\mathcal{A} \vDash \varphi(\sigma a_1, ..., \sigma a_n)$, and it follows that $\hat{\sigma}b$ is in $||\mathtt{resp}(\varepsilon, \mathcal{A})||$, as desired.

For the second claim, the verification when ε is *not* an `invent`-expression is routine, so we assume again that ε is of the form

$$(X \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\})$$

Thus, assuming $\overline{b} = (b_1, ..., b_r)$, for each $j = 1, ..., r$, there are $i_j \in \{1, ..., k - n\}$ and $(a_1^j, ..., a_n^j) \in \{\overline{x}|\varphi\}(\mathcal{A})$ such that $b_j = \langle i_j, a_1^j, ..., a_n^j\rangle$. Note that every $\overline{b}'$ in the orbit of $\overline{b}$ under $Aut(\mathtt{resp}(\varepsilon, \mathcal{A}))$ is of the form $\overline{b}' = (b_1', ..., b_r')$, where for each $j = 1, ..., r$, there are $l_j \in \{1, ..., k - n\}$ and $(c_1^j, ..., c_n^j) \in \{\overline{x}|\varphi\}(\mathcal{A})$ such that $b_j' = \langle l_j, c_1^j, ..., c_n^j\rangle$. (Then for $\sigma \in Aut(\mathcal{A})$, $\hat{\sigma}\overline{b} = \overline{b}'$ just in case $l_j = i_j$ for each $j = 1, ..., r$ and $\sigma(b_i^j) = \sigma(c_i^j)$ for all $j = 1, ..., r$ and $i = 1, ..., n$.) If $\sigma \in Aut(\mathtt{resp}(\varepsilon, \mathcal{A}))$ takes $\overline{b}$ to $\overline{b}'$, then as σ preserves $\mathtt{invent}_k\{\overline{x}|\varphi\}(\mathcal{A})$, we have $l_j = i_j$ for each $j = 1, ..., r$. Moreover, for each j,

$$\sigma(a_1^j, ..., a_n^j, \langle i_j, a_1^j, ..., a_n^j\rangle) = (c_1^j, ..., c_n^j, \langle l_j, c_1^j, ..., c_n^j\rangle)$$

Thus, for each automorphism $\sigma \in Aut(\mathtt{resp}(\varepsilon, \mathcal{A})$, there is an automorphism τ of the structure $\mathcal{A}[\, X/\{\overline{x}|\varphi\}(\mathcal{A})\,]$ such that $\sigma = \hat{\tau}$, and this suffices to prove the claim. (When $\overline{b}$ consists of both elements of A and invented elements, the proof is not significantly different, so we leave the demonstration here.) □

Proposition 3.3. *Let $\mathcal{A}$ be a situated ρ_1-structure, and let P be an* `invent`*-program over ρ_1 such that $\mathtt{resp}(P, \mathcal{A})$ is defined. Then, if $0 < r < \omega$ and $\overline{b} \in ||\mathtt{resp}(P, \mathcal{A})||^r$, then $\overline{b}^{Aut(\mathcal{A})} \subseteq ||\mathtt{resp}(P, \mathcal{A})||^r$ and $\overline{b}^{Aut(\mathcal{A})}$ is a union of orbits of $Aut(\mathtt{resp}(P, \mathcal{A}))$ on r-tuples.*

Proof. The demonstration is a straightforward inductive argument using the fact that evaluating $\mathtt{resp}(P, \mathcal{A})$ is identical to evaluating a certain sequence of basic expressions. $\square$

The following theorem is also from [8].

Theorem 3.4. *Every HL-transformation is computable by an* `invent`*-program up to isomorphism over universes of input structures.*

3.2.3 Some easy program normalizations

Let $\rho_1 \subset \rho_1'$, where ρ_1' is finite, and assume $V \in \rho_1' \setminus \rho_1$ is a unary relation variable. Let P be a program over ρ_1 and P' a program over ρ_1'. Let K be a set of situated ρ_1-structures which is closed under isomorphisms. For any situated ρ_1-structure $\mathcal{A}$, there is a "trivial" ρ_1'-expansion $\mathcal{A}'$ of $\mathcal{A}$ in which $Y^{\mathcal{A}'} = \emptyset$ for all $Y \in \rho_1' \setminus \rho_1$. We say that P' *captures* P *over* K if for every $\mathcal{A} \in K$, if $\mathtt{resp}(P, \mathcal{A})$ is defined, then $\mathtt{resp}(P', \mathcal{A}')$ is defined and, setting $\mathcal{B} = \mathtt{resp}(P', \mathcal{A}')$,

$$\left\langle V^{\mathcal{B}}, \left(X^{\mathcal{B}} \cap (V^{\mathcal{B}})^{ari(X)} : X \in \rho_1\right)\right\rangle \cong \mathtt{resp}(P, \mathcal{A})$$

We say that a basic expression $\varepsilon = (X \leftarrow \{\overline{x}|\varphi\})$ or $\varepsilon = (X \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\})$ is *flat* just in case $\varphi(x_1, ..., x_n)$ is of one of the the following forms:

1. An equality type, $\tau(x_1, ..., x_n)$;
2. A literal $\pm Y(x_1, ..., x_n)$, where $Y \in \rho_1$;
3. A conjunction or disjunction:

   $$Y_1(x_{i_1}, ..., x_{i_r}) \wedge Y_2(x_{j_1}, ..., x_{j_s}) \text{ or } Y_1(x_{i_1}, ..., x_{i_r}) \vee Y_2(x_{j_1}, ..., x_{j_s})$$

   where $\{x_{i_u}, x_{j_v}, u \in [r], v \in [s]\} = \{x_1, ..., x_n\}$.
4. An existential $\exists x_{n+1} ... x_{n+t}\, Y(x_1, ..., x_n, x_{n+1}, ..., x_{n+t})$. (In this case, we allow $x_{n+i} \in \{x_1, ..., x_n\}$, $i = 1, ..., t$.

The following propositions are folklore:

Proposition 3.5. *Let P be a program over ρ_1. Then there are $\rho_1' \supset \rho_1$, a program P' over ρ_1' and a number $d < \omega$ such that*

1. *P' captures P over the set of all situated ρ_1-structures.*
2. $\mathtt{cpx}(P', \mathcal{A}') \leq (2 + \mathtt{cpx}(P, \mathcal{A}))^d$ *whenever* $\mathtt{resp}(P, \mathcal{A})$ *is defined.*
3. *P' is of the form $Q_0; (\textbf{\textit{while}}\, \psi\, \textbf{\textit{do}}\, Q)$; where Q is loop-free and Q_0 is a sequence of flat basic expressions.*

Proposition 3.6. *Let Q be a loop-free program over ρ_1. Then there are $\rho_1' \supset \rho_1$, a program Q' over ρ_1' and a number $c < \omega$ such that*

1. *P' captures P over the set of all situated ρ_1-structures.*
2. $\mathtt{cpx}(Q', \mathcal{A}') \leq c \cdot (2 + \mathtt{cpx}(Q, \mathcal{A}))$ *for all situated ρ_1-structures $\mathcal{A}$.*
3. *Q' is of the form*

$$(\textbf{\textit{if}}\, \psi_1\, \textbf{\textit{then}}\, R_1\, \textbf{\textit{else}}\, id); ...; (\textbf{\textit{if}}\, \psi_m\, \textbf{\textit{then}}\, R_m\, \textbf{\textit{else}}\, id)$$

*where $\emptyset \vdash \bigvee_{i=1}^m \psi_i$, $\psi_i \wedge \psi_j \vdash$ **false** whenever $i \neq j$, and for each $i = 1, ..., m$, R_i is a sequence of flat basic expressions.*

3.2.4 Essentially inflationary programs and constructible theories

Throughout this section, we consider *fully normalized* `invent`-programs P over a signature ρ_1. That is, P is of the form $P_{pre}; P_{loop}$ – and more precisely of the form:

$$P_{pre}; (\ \textbf{while}\, \varphi_{loop}\, \textbf{do}\, (\textbf{if}\, \psi_1\, \textbf{then}\, R_1\, \textbf{else}\, \text{id}); ...; (\textbf{if}\, \psi_m\, \textbf{then}\, R_m\, \textbf{else}\, \text{id})\ )$$

satisfying the following conditions:

1. P_{pre} is a sequence of basic expressions.
2. $\emptyset \vdash \bigvee_i \psi_i$ and if $i \neq j$, then $\psi_i \wedge \psi_j \vdash$ **false**
3. For each $i = 1, ..., m$, R_i is a sequence of flat basic expressions.

We will also desire that P_{loop} is *essentially inflationary* in a sense to be made precise next. A relation variable $X \in \rho_1$ is called *private* (with respect to P) if (a) it does not occur in any of $\psi_1, ..., \psi_m$ or φ_{loop}, and (b) for each $i = 1, ..., m$, X does not occur in the body of any basic expression of R_i before it has occurred as the head of a basic expression in R_i. Naturally enough, X is called *public* with respect to P if it is not private, and we let $pub(P)$ denote the set of public relation variables of P. We say that P is *essentially inflationary* if for each $i = 1, ..., m$, for each $X \in pub(P)$, every basic expression in R_i with X as the head is the form $(X \leftarrow \{\overline{x} | X(\overline{x}) \vee X_1(\overline{x})\})$ – that is, R_i is *explicitly inflationary* with respect to X.

We are at last in a position to specify our notion of a *(fully) coherent solution* of the model-building problem for a k-variable theory T. Let P be a fully normalized `invent`-program over a signature $\rho_1 \supseteq \rho^G$ (with respect to T); then P is a (fully) coherent solution of the model-building problem for T just in case:

1. P is a weakly coherent solution of the model-building problem for T;
2. P is essentially inflationary;

3. $\rho^G \subseteq pub(P)$.

As before, T is *(fully) constructible* if it admits a fully coherent solution for the model-building problem. In the next chapter, we will summarize how to show that any *polynomial-time* weakly coherent solution for the model-building is necessarily fully coherent, and in chapter 8, we will also show how to extract an independence relation on $\mathbf{fin}[T^G]$ from the hypothesis of full constructibility.

3.3 Fixed-point logics over HL and HL$^{(t)}$

3.3.1 Some background on fixed-point logics

Let ρ be a finite relational signature. We define the fixed-point logic $FP[\rho]$, and extension of the first-order language of ρ, as follows:

1. Every first-order formula is in $FP[\rho]$, and $FP[\rho]$ is closed under the boolean connectives and first-order quantification.

2. Suppose $\varphi = \varphi(x_1, ..., x_n, \overline{y}; X^{(n)}, Y_1, ..., Y_m)$ is a formula of $FP[\rho \cup \{X, Y_1, ..., Y_m\}]$. Then
$$[\mathtt{fp} : X, (x_1, ..., x_n) : \varphi](v_1, ..., v_n, \overline{y}; \overline{Y})$$
is in $FP[\rho \cup \{Y_1, ..., Y_m\}]$

For the most part, we will only deploy the syntax of fixed-point logic in the analysis of essentially inflationary programs, but for completeness, we will provide the so-called partial (`pfp`) semantics in addition to the inflationary (`ifp`) semantics of the fixed-point logic. Let V be a nonempty set, and let $F : \mathscr{P}(V) \to \mathscr{P}(V)$. We then define,

$$\begin{aligned}
F^0 &= \emptyset \\
F^{\alpha+1} &= F^\alpha \cup F(F^\alpha) \\
F^\lambda &= \bigcup_{\alpha<\lambda} F^\alpha \text{ (if } \lambda \text{ is a limit ordinal)} \\
\mathtt{ifp}(F) &= \bigcup_{\alpha<2^{|S|}} F^\alpha
\end{aligned}$$

The `ifp`-semantics of $FP[\rho]$ are then given as follows:

1. The semantics of first-order formulas is unchanged, as are the conditions for satisfaction of conjunctions, negations and first-order quantification.

2. Suppose $\varphi = \varphi(x_1, ..., x_n, \overline{y}; X^{(n)}, Y_1, ..., Y_m)$ is a formula of the "larger" fixed-point logic $FP[\rho \cup \{X, Y_1, ..., Y_m\}]$. Let $\mathcal{M}' = (\mathcal{M}, R_1, ..., R_m)$ be a $\rho \cup \{Y_1, ..., Y_m\}$-structure, where $\mathcal{M}$ is a ρ-structure, and let $\overline{b} \in M^{\overline{y}}$. Define $F_{\varphi, \bar{b}, \overline{R}} : \mathscr{P}(M^n) \to \mathscr{P}(M^n)$ by

   $$F_{\varphi, \bar{b}, \overline{R}}(S) = \left\{\overline{a} \in M^n : \mathcal{M}' \vDash \varphi(\overline{a}, \overline{b}; S, \overline{R})\right\}$$

   Then

   $$\mathcal{M}' \vDash [\mathtt{fp} : X, (x_1, ..., x_n) : \varphi](\overline{a}, \overline{b}) \;\Leftrightarrow\; \overline{a} \in \mathtt{ifp}(F_{\varphi, \bar{b}, \overline{R}})$$

We now move to the $\mathtt{pfp}$-interpretation of $FP[\rho]$. Again, let V be a nonempty set, and let $F : \mathscr{P}(V) \to \mathscr{P}(V)$. We then define,

$$\begin{aligned} F_p^0 &= \emptyset \\ F_p^{m+1} &= F(F_p^m) \\ \mathtt{pfp}(F) &= \begin{cases} F_p^n & \text{if } n = \min\left\{m : F_p^{m+1} = F_p^m\right\} \text{ exists} \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

The rest of the interpretation is unchanged from the $\mathtt{ifp}$-interpretation.

Constructions using simultaneous fixed-points are also of interest, although it can be proved that simultaneous fixed-point constructions do not contribute any additional expressive power. Suppose

$$\Psi = \{\psi_1(\overline{x}_1, ..., \overline{x}_m; X_1, ..., X_m), ..., \psi_m(\overline{x}_1, ..., \overline{x}_m; X_1, ..., X_m)\}$$

are formulas of $FP[\rho \cup \{X_1, ..., X_m\}]$ were $ari(X_i) = r_i$ for each $i = 1, ..., m$. We have an additional formula construction,

$$[\mathtt{fp} : X_i, x_i^1, ..., x_i^{r_i} : \Psi](v_1, ..., v_{r_i})$$

The associated operator for a ρ-structure $\mathcal{M}$,

$$F_\Psi : \mathscr{P}(M^{r_1} \times \cdots \times M^{r_m}) \longrightarrow \mathscr{P}(M^{r_1} \times \cdots \times M^{r_m})$$

is given by

$$(R_1, ..., R_m) \mapsto \left\{(\overline{a}_1, ..., \overline{a}_m) : \bigwedge_{i=1}^{m} \mathcal{M} \vDash \psi_i(\overline{a}_1, ..., \overline{a}_m; \overline{R})\right\}$$

The $\mathtt{ifp}$- and $\mathtt{pfp}$-semantics are then nearly identical to those for single-formula fixed-points – computing the simultaneous fixed-point of the system (if it exists) and simply extracting the desired relation.

3.3.2 Fixed-point logic over $\mathrm{HL}^{(t)}$

Recall that the signature of HL^t is $\tau^t_{\mathrm{HL}} = \{J_1, ..., J_t, 1, ..., t\}$, where $1, ..., t$ are constant symbols and for each $r = 1, ..., t$,

$$J_r^{\mathrm{HL}^{(t)}} = \left\{(a_1, ..., a_r, \langle a_1, ..., a_r\rangle) : a_1, ..., a_r \in \mathrm{HL}^{(t)}\right\}$$

If ρ is a finite relational signature (disjoint from τ^t_{HL}), then we define $\rho + \tau^t_{\mathrm{HL}}$ to be the signature $\rho \,\dot\cup\, \tau^t_{\mathrm{HL}} \,\dot\cup\, \{D\}$.

Lemma 3.7. *Let ε be a basic expression over ρ of arity $k \leq t$. Then there is a first-order formula $\theta_\varepsilon(x_1, ..., x_n)$ of $\rho + \tau^t_{HL}$ such that for all ρ-structures $\mathcal{A}$ with $A \subseteq HL^{(t)}$,*

$$(a_1, ..., a_k) \in \varepsilon(\mathcal{A}) \;\Leftrightarrow\; (HL^{(t)}, \mathcal{A}) \vDash \theta_\varepsilon(a_1, ..., a_k)$$

for all $a_1, ..., a_k \in HL^{(t)}$.

Proof. First, if $\varepsilon = (X \leftarrow \{\overline{x}|\varphi\})$, where $\varphi(x_1, ..., x_k)$ is in $FO[\rho]$, then $\theta_\varepsilon(\overline{x})$ is simply φ^D. Now, suppose $\varepsilon = (X \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\})$ where $\varphi(x_1, ..., x_n)$ is in $FO[\rho]$ and $n < k$. Then, $\theta_\varepsilon(x_1, ..., x_k)$ is the following formula:

$$\varphi^D(x_1, ..., x_n) \wedge \bigwedge_{i=n+1}^{k} J_{n+1}(i - n, x_1, ..., x_n, x_i) \tag{3.1}$$

The verification is routine, and we omit it. □

Proposition 3.8. *Let P be an* `invent`*-program over ρ, where $ari(X) \leq t$ for all $X \in \rho$. Then there is a set of first-order formulas $\left\{\theta^X_P(x_1, ..., x_n)\right\}_{X\in\rho}$ of $\rho + \tau^t_{HL}$ such that for all ρ-structures $\mathcal{A}$ with $A \subseteq HL^t$ and every $X^{(n)} \in \rho$,*

$$(a_1, ..., a_n) \in X^{\mathtt{resp}(P,\mathcal{A})} \;\Leftrightarrow\; (HL^t, \mathcal{A}) \vDash \theta^X_P(a_1, ..., a_n)$$

for all $a_1, ..., a_n \in HL^t$. Moreover, if P is fully normalized and essentially inflationary, then we may use the `ifp`*-semantics in the evaluation of the formula.*

Proof. We deal with the several construction rules for `invent`-programs.

<u>P is a basic expression with head $X^{(r)}$</u>: For $Y^{(n)} \in \rho\backslash\{X\}$, $\theta^Y_P(x_1, ..., x_n)$ is just $Y(x_1, ..., x_n)$, and $\theta^X_P(x_1, ..., x_r)$ is $\theta_P(x_1, ..., x_r)$ as in the previous lemma.

<u>$P = P_1; P_2$</u>: Let $X \in \rho$. Obtain θ^X_P be replacing occurrence $Y(v_1, ..., v_n)$ of Y in $\theta^X_{P_2}$ by $\theta^Y_{P_1}(v_1, ..., v_n)$ for each $Y^{(n)} \in \rho$.

$P = (\textbf{if}\,\varphi\,\textbf{then}\,P_1\,\textbf{else}\,P_2)$: For each $X^{(n)} \in \rho$, let $\theta_P^X(x_1, ..., x_n)$ be

$$(\varphi^D \wedge \theta_{P_1}^X(\overline{x})) \vee ((\neg\varphi)^D \wedge \theta_{P_2}^X(\overline{x}))$$

$P = (\textbf{while}\,\varphi\,\textbf{do}\,P_1)$: For each $X^{(n)} \in \rho$, let $\psi_P^X(x_1, ..., x_n)$ be

$$(\varphi^D \wedge \theta_{P_1}^X(\overline{x})) \vee ((\neg\varphi)^D \wedge X(\overline{x}))$$

$$[\mathtt{fp} : X, x_1, ..., x_n : \Psi_P](x_1, ..., x_n)$$

Again, the verification is straightforward, so we omit it. □

To obtain a converse of proposition 3.8, we specialize the full fixed-point logic (syntactically speaking) to a natural fragment $FP^*[\rho]$:

1. Every first-order formula of the form φ^D or of that in lemma 3.7 is an $FP^*[\rho]$-formula, and $FP^*[\rho]$ is closed under boolean connectives and first-order quantification.

2. The fixed-point formula construction is unchanged.

The proof of the following proposition is not hard, but quite long and tedious – we omit the demonstration in the interest of not boring the reader into an early grave.

Proposition 3.9. *Let $\Psi = \left\{\psi^X(x_1, ..., x_r) : X^{(r)} \in \rho\right\}$ be a family of $FP^*[\rho]$-formulas, possibly involving simultaneous fixed-point constructions.*

1. *(Evaluation under `pfp`-semantics) There is a fully normalized `invent`-program P_Ψ such that for every initial ρ-structure $\mathcal{A}$, $\mathtt{resp}(P_\Psi, \mathcal{A})$ is defined if and only if*

$$\left\{\overline{a} \in (HL^t)^r : (HL^t, \mathcal{A}) \vDash [\mathtt{pfp} : X, \overline{x} : \Psi](\overline{a})\right\}$$

is finite for each $X^{(r)} \in \rho$. Moreover, if $\mathtt{resp}(P_\Psi, \mathcal{A})$ is defined, then

$$X^{\mathtt{resp}(P_\Psi, \mathcal{A})} = \left\{\overline{a} \in (HL^t)^r : (HL^t, \mathcal{A}) \vDash [\mathtt{pfp} : X, \overline{x} : \Psi](\overline{a})\right\}$$

for each $X \in \rho$.

2. *(Evaluation under `ifp`-semantics) There is a fully normalized essentially inflationary `invent`-program Q_Ψ satisfying the same conditions that P_Ψ did in the previous statement.*

Chapter 4

Characterizations of efficient transformations

In this chapter, we will sketch the proofs of some program normalizations that apply specifically to efficient programs, showing in particular that every efficient program is captured by an efficient essentially inflationary program. Once again, the analysis is only mildly different from that of [1], and our contribution is simply to extend the technique to accommodate object-creation. Before embarking on that project, however, we pause to show how our model of computation fits in with more "standard" machine-based ones; hopefully, this demonstration will satisfy the reader that our analysis is reasonably robust to details of the model of computation. Through this chapter's analyses it will become clear that in order to characterize efficient model-building, it suffices to consider essentially inflationary programs. Using this fact, we will see, in chapter 8, that if T is efficiently constructible, then $\mathbf{fin}[T^G]$ is rosy.

4.1 Relational Turing machines

In this section, we introduce the notion of a relational Turing machine – a standard Turing machine interacting with a structure exclusively through first-order definable transformations of the structure. (The idea of such a machine seems to go back to [1] and related papers by the same authors, where it is thought of as a database query language, like SQL, embedded in a Turing-complete programming language, like C.) Formally, a relational Turing machine (RTM) is a tuple $M = (\Sigma, Q, q_0, F, \rho_1, S, B, \delta)$ such that:

1. Σ is a finite alphabet. Q is a finite set of states; $q_0 \in Q$ is the initial state; and $F \subseteq Q$ is a set of final or halting states.

2. ρ_1 is a finite relational signature; S is a finite set of $FO[\rho_1]$-sentences; B is a finite set of basic expressions over ρ_1.

3. $\delta : Q \times (\Sigma \dot{\cup} \{\Box\}) \times \mathscr{P}(S) \longrightarrow Q \times \Sigma \times \{l, r\} \times B$ is the transition function.

A configuration of M is a pair $(w, \mathcal{A})$, where w is a word of the regular language $\Box\Sigma^*(Q \times (\Sigma \cup \{\Box\}))\Sigma^*\Box$ and $\mathcal{A}$ is a ρ_1-structure with $A \subseteq$ HL. In particular, w takes one of the following forms:

$$\begin{array}{rl} \Box w'x(q,z)yw''\Box & \text{where } x, y \in \Sigma,\ q \in Q, \text{ and } w', w'' \in \Sigma^* \\ \Box(q,z)yw''\Box & \text{where } y \in \Sigma,\ q \in Q, \text{ and } w'' \in \Sigma^* \\ \Box w'x(q,z)\Box & \text{where } x \in \Sigma,\ q \in Q, \text{ and } w' \in \Sigma^* \\ \Box(q,z)\Box & \text{where } q \in Q \end{array}$$

and in each case $z \in \Sigma \cup \{\Box\}$. The machine's read/write head is understood to be scanning the symbol $z \in \Sigma\dot{\cup}\{\Box\}$ paired with the state symbol. If $q \in F$, then the machine has halted, so there is no succeeding configuration. Assuming, then, that $q \notin F$, we set $S_{\mathcal{A}} = \{\psi \in S : \mathcal{A} \vDash \psi\}$, and consider $(z_1, q_1, d, \varepsilon) = \delta(q, z, S_{\mathcal{A}})$ where $d \in \{l, r\}$ and $\varepsilon \in B$ is a basic expression. We deal with transitions from each of the four configuration types in turn:

1. $w = \Box w'x(q,z)yw''\Box$:

   If $d = l$, then the next configuration is $(\Box w'(q_1, x)z_1yw''\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$, and if $d = r$, then it is $(\Box w'z_1(q_1, y)w''\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$.

2. $w = \Box(q,z)yw''\Box$:

   If $d = l$, then the next configuration is $(\Box(q_1, z_1)yw''\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$, and if $d = r$, then it is $(\Box z_1(q_1, y)w''\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$. (The machine's tape is infinite only to the right)

3. $w = \Box w'x(q,z)\Box$:

   If $d = l$, then the next configuration is $(\Box w'(q_1, x)z_1yw''\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$, and if $d = r$, then it is $(\Box z_1(q_1, \Box)\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$.

4. $w = \Box(q,z)\Box$:

   If $d = l$, then the next configuration is $(\Box(q_1, z_1)\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$, and if $d = r$, then it is $(\Box z_1(q_1, \Box)\Box, \texttt{resp}(\varepsilon, \mathcal{A}))$.

As a convention, the initial configuration of an RTM-computation on a structure $\mathcal{A}$ will always be $(\Box(q_0, \Box)\Box, \mathcal{A})$ – that is, the tape is initially empty. The output of an RTM-computation *output*$(M, \mathcal{A})$ is simply the structure whose universe comprises all elements that occur in a relation register at any point of the computation and whose relations are those stored in the relation registers upon termination. The time-complexity of an RTM computation, $\texttt{cpx}_{time}(M, \mathcal{A})$, is measured simply as the number of steps in the computation until it halts. The space-complexity, $\texttt{cpx}_{space}(M, \mathcal{A})$, is measured as the sum of (a) the number of non-$\Box$-marked cells left on the tape at the end of the computation, (b) n^k where k is the maximum arity of a relation symbol of M and n is the cardinality of $||output(M, \mathcal{A})||$ For a function $f : \mathbb{N} \to \mathbb{N}$, an HL-transformation Q is in REL-TIME$(f(n))$ just in case there is an RTM M such that $\mathcal{A} \mapsto output(M, \mathcal{A})$ is identical to Q (up to isomorphism over A) on $dom(Q)$ and there is a number c such that for almost all $\mathcal{A} \in dom(Q)$, $\mathtt{cpx}_{time}(M, \mathcal{A}) \leq c \cdot f(|A|)$. REL-SPACE$(f(n))$ is defined similarly. We then define the analogues of the classical complexity classes, PTIME and PSPACE, as follows:

$$\text{REL-PTIME} = \bigcup_{t<\omega} \text{REL-TIME}(n^t)$$
$$\text{REL-PSPACE} = \bigcup_{t<\omega} \text{REL-SPACE}(n^t)$$

Note that due to the presence of `invent`-operations, these definitions of REL-PTIME and REL-PSPACE are not trivially identical to the definitions posited in [1]; it *can* be shown, however, that our definitions and those of [1] coincide for decision problems. In [1], it is also proved that REL-PTIME coincides with fixed-point logic under `ifp`-semantics, and REL-PSPACE coincides with fixed-point logic under `pfp`-semantics (under their definitions). Further, it is proved that REL-PSPACE contains exactly those decision problems computable by unrestricted **while**-programs (our `invent`-programs, omitting the `invent`-operation), and more saliently, REL-PTIME contains exactly those decision problems computable by inflationary **while**-programs. In the remainder of this chapter, we will sketch how to adapt the methods of [1] to prove similar equivalences for `invent`-programs and HL-transformations.

4.1.1 Equivalences over ordered initial structures

In this section, we assume that our finite relational signature ρ_1 contains a binary relation symbol $<$. An ordered initial ρ_1-structure, then, is an initial ρ_1-structure $\mathcal{A}$ such that $<^{\mathcal{A}}$ is a linear order of its universe A. The arguments for the following propositions are quite standard (requiring only very minor modifications) – templates for all of them can be found in [19] or [11]. All of them boil down to more or less elaborate encoding excercises.

Proposition 4.1. *Let K be a set of ordered initial ρ_1-structures, closed under isomorphism, and let M be an RTM over a signature $\rho_1' \supseteq \rho_1$. Suppose that M has polynomial running time over K – that is, for some $d < \omega$,*

$$\mathtt{cpx}_{time}(M, \mathcal{A}) \leq (2 + |A|)^d$$

for all $\mathcal{A} \in K$. Then there is a family $\Psi = \{\psi^X : X \in \rho_1'\}$ of first-order formulas of $\rho_1' + \tau_{HL}^t$ for some $0 < t < \omega$ such that for all $\mathcal{A} \in K$ and every $X^{(r)} \in \rho_1'$,

$$X^{output(M,\mathcal{A})} = \{(a_1, ..., a_r) \in (HL^t)^r : (HL^t, \mathcal{A}) \models [\mathtt{fp} : X, \overline{x} : \Psi](\overline{a})\}$$

evaluating the simultaneous fixed-point under the `ifp`*-semantics.*

The following corollary is, then, an immediate consequence of the preceding proposition together with proposition 3.9 of chapter 3.

Corollary 4.2. *Let K be a set of ordered initial ρ_1-structures, closed under isomorphism, and let M be an RTM over a signature $\rho_1' \supseteq \rho_1$. Suppose that M has polynomial running time over K. Then there is a fully normalized essentially inflationary* `invent`*-program P_M over $\rho_1'' \supseteq \rho_1'$ and a number $c < \omega$ such that for all $\mathcal{A}$*

$$\mathtt{cpx}(P_M, \mathcal{A}) \leq c \cdot (2 + \mathtt{cpx}_{time}(M, \mathcal{A}))$$

and for every $X \in \rho_1'$, $X^{output(M,\mathcal{A})} = X^{\mathtt{resp}(P_M,\mathcal{A})}$ and $X \in pub(P_M)$.

As one might expect, the RTM model of computation is universal for "relational" models of computation – those which interact with structures only through first-order definable transformations. In particular, we have:

Lemma 4.3. *Let K be a set of (not necessarily ordered) initial ρ_1-structures, closed under isomorphism, and let P be an* `invent`*-program over a signature $\rho_1' \supseteq \rho_1$. Then there is an RTM M_P over ρ_1 and a number $c' < \omega$ such that for all $\mathcal{A} \in K$,*

$$\mathtt{cpx}_{time}(M_P, \mathcal{A}) \leq c' \cdot (2 + \mathtt{cpx}(P, \mathcal{A}))$$

and $X^{output(M_P,\mathcal{A})} = X^{\mathtt{resp}(P,\mathcal{A})}$ for every $X \in \rho_1$.

To conclude this section, we note that the dependencies on the linear order in these facts is not as strict as it may appear. More precisely, if $<$ and E are $2k$-ary relation symbols in ρ_1 such that for all $\mathcal{A} \in K$, (i) $E^{\mathcal{A}}$ is an equivalence relation on A^k, (ii) $<^{\mathcal{A}}$ is a linear order of $A^k/E^{\mathcal{A}}$, and (iii) the action of the program or machine "respects E-classes" in a certain sense (see the next section), then we obtain essentially the same results provided that we simply measure complexity as a function of $|A^k/E|$ rather than $|A|$. If the class K admits an ensemble $(E, <)$ of this kind (explicitly or definably), then we say that K is *pseudo-ordered* by $(E, <)$, and $(E, <)$ is a k-ary pseudo-order for K.

4.2 Reduction to pseudo-ordered structures

Before we can proceed, there are some necessary pieces of machinery to introduce first. To smooth the discussion, we will assume that our finite relational signature ρ_1 contains only relation symbols of arity exactly k for some positive integer k. Let Φ be a set of formulas of the language of ρ_1 of arity $\leq k$. Φ may be infinite, but in general, we will assume that $\Phi_{\rho_1} \subseteq \Phi$, where

$$\Phi_{\rho_1} = \{R(x_1, ..., x_k) : R \in \rho_1\} \cup \left\{ \bigwedge_{i<j} (x_i = x_j)^{\sigma(i,j)} : \sigma : [k] \times [k] \to \{0, 1\} \right\}$$

for some fixed set of pairwise distinct variables $\{x_1, ..., x_k\}$. We say, then, that Φ is *acceptable*. Let $\mathcal{A}$ be a ρ_1-structure, and let $(E, <)$ be a k-ary pseudo-order of $\mathcal{A}$. We say that $(E, <)$ is a Φ*-respecting pseudo-order of* $\mathcal{A}$ just in case for each $\varphi(x_1, ..., x_k) \in \Phi$, there is a set $\{C_1, ..., C_m\}$ of E-classes such that $\{\overline{x}|\varphi\}(\mathcal{A}) = \bigcup_{i=1}^{m} C_i$. Assuming Φ is finite, we define another, larger set of formulas Φ^* as follows:

1. Φ_0 is the set of formulas of the form

$$\varphi(x_1, ..., x_n) \wedge \bigwedge_{i=n+1}^{k} x_i = x_n$$

   for $\varphi(x_1, ..., x_n) \in \Phi$.

2. Given Φ_m, we immediately put $\Phi_m \subseteq \Phi_{m+1}$.

   Suppose $\varphi(\overline{x}) \in \Phi_0$ and $R_1(\overline{v}_1), ..., R_s(\overline{v}_s)$ is an enumeration of the atomic subformulas of φ, with repetitiions . Let $\psi_1(\overline{x}), ..., \psi_s(\overline{x}) \in \Phi_m$, and let $\varphi[\psi_1, ..., \psi_s]$ be the formula obtain by replacing $R_i(\overline{v}_i)$ by $\psi_i(\overline{v}_i)$ for each $i = 1, ..., s$. Then $\varphi[\psi_1, ..., \psi_s]$ is in Φ_{m+1}. (Note the $\varphi[\psi_1, ..., \psi_s]$ and φ have the same set of free variables, and in particular, every formula in Φ_{m+1} is k-ary.)

3. $\Phi^* = \bigcup_{m<\omega} \Phi_m$.

$(E, <)$ is a *finest* Φ^*-respecting pseudo-order of $\mathcal{A}$ if for all $\overline{a}, \overline{b} \in A^k$, $(\overline{a}, \overline{b}) \in E$ implies that $\mathcal{A} \models \varphi(\overline{a})$ iff $\mathcal{A} \models \varphi(\overline{b})$ for all $\varphi(x_1, ..., x_k) \in \Phi^*$.

Theorem 4.4. *[1] Let Φ be an acceptable finite set of ρ_1-formulas. There is an inflationary* ***while****-program Q_Φ over $\rho_1 \cup \{E^{(2k)}, <^{(2k)}\}$ such that the following holds:*

Suppose $\mathcal{A}$ is a ρ_1-structure and $(E_0, <_0)$ is Φ_{ρ_1}-respecting pseudo-order of $\mathcal{A}$. Then, if $(\mathcal{A}, E_1, <_1) = \mathtt{resp}(Q_\Phi, (\mathcal{A}, E_0, <_0))$, then $(E_1, <_1)$ is a finest Φ^*-respecting pseudo-order of $\mathcal{A}$ such that

1. $(\overline{a}, \overline{b}) \in E_1 \Rightarrow (\overline{a}, \overline{b}) \in E_0$,

2. $\overline{a}/E_0 <_0 \overline{b}/E_0 \Rightarrow \overline{a}/E_1 <_1 \overline{b}/E_1$,

for all $\overline{a}, \overline{b} \in A^k$. Moreover, $\mathtt{cpx}(Q_\Phi, (\mathcal{A}, E_0, <_0)) \in |A|^{k \cdot O(1)}$.

For the next few lemmas, it will be convenient to define

$$\mathtt{invent}_k(\overline{a}) = (a_1, ..., a_n, 1\frown\overline{a}, ..., (k-n)\frown\overline{a})$$

$(pad_k(\overline{a})$ was defined in the introduction.)

Lemma 4.5. *Let Φ be an acceptable finite set of ρ_1-formulas. Suppose $(E_0, <_0)$ is a finest Φ^*-respecting pseudo-order of a situated ρ_1-structure $\mathcal{A}$.*

Let $\varphi(x_1, ..., x_n) \in \Phi$ and $\varepsilon = (R \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\})$, and let $\overline{a}, \overline{b} \in \{\overline{x}|\varphi\}(\mathcal{A})$ and $\overline{c}, \overline{d} \in A^k$. Suppose $(\overline{c}, \overline{d}) \in E_0$ and $(pad_k\overline{a}, pad_k\overline{b}) \in E_0$, and suppose $X, Y \subseteq [k]$ is a partition of $[k]$. Then

$$\mathtt{resp}(\varepsilon, \mathcal{A}) \vDash \psi(\ (\overline{c}{\restriction}X)^\frown(\mathtt{invent}_k(\overline{a}){\restriction}Y)\ )$$

if and only if

$$\mathtt{resp}(\varepsilon, \mathcal{A}) \vDash \psi(\ (\overline{d}{\restriction}X)^\frown(\mathtt{invent}_k(\overline{b}){\restriction}Y)\ )$$

for all $\psi(x_1, ..., x_k) \in \Phi^*$.

Let $<^k_{\rho_1}$ be a linear order of the set $S^{\mathrm{qf}}_k(\rho_1)$ of complete quantifier-free k-types of the language of ρ_1. If $\mathcal{A}$ is a ρ_1-structure, there is a natural Φ_{ρ_1}-respecting pseudo-order $(E^k_{\rho_1}, <^k_{\rho_1})$ of $\mathcal{A}$. By theorem 4.4, there is a canonical finest Φ^*-respecting pseudo-order $(E^{\mathcal{A}}, \prec^{\mathcal{A}})$ of $\mathcal{A}$ modulo $(E^k_{\rho_1}, <^k_{\rho_1})$ and the program Q_Φ, and this is computable in relational polynomial-time just because the `invent`-operator is not in play.

$\star\star\star$ For the rest of this section, we assume that P is a fully normalized program over ρ_1, and we assume that there is a fixed set $\{x_1, ..., x_k\}$ of first-order variables such that all control sentences and formulas appearing in P take their variables from this set.

Let Φ_P be the set of formulas that appear in P, whether as control or in basic expressions, together with every formula of Φ_{ρ_1}; obviously, this set is finite. Let ρ_1° be the following signature:

$$\left\{\prec^{(2)}\right\} \cup \left\{E^{(2)}_X : \emptyset \neq X \subseteq [k]\right\} \cup \left\{F^{(2)}_\sigma : \sigma \in Sym\,[k]\right\} \cup \left\{V_q : q \in S^{\mathrm{qf}}_k(\rho_1)\right\}$$

Given a ρ_1-structure, we define canonically a ρ_1°-structure $\mathcal{A}/P$ as follows:

- $||\mathcal{A}/P|| = A^k/E^{\mathcal{A}}$, $\prec^{\mathcal{A}/P} = \prec^{\mathcal{A}}$
- $E^{\mathcal{A}/P}_X = \left\{(\overline{a}/E^{\mathcal{A}}, \overline{b}/E^{\mathcal{A}}) : \bigwedge_{i \in X} a_i = b_i\right\}$
- $F^{\mathcal{A}/P}_\sigma = \left\{(\overline{a}/E^{\mathcal{A}}, \overline{b}/E^{\mathcal{A}}) : \overline{b} = (a_{\sigma(1)}, ..., a_{\sigma(k)})\right\}$
- $V^{\mathcal{A}/P}_q = \left\{\overline{a}/E^{\mathcal{A}} : \mathcal{A} \vDash q(\overline{a})\right\}$

There is also a natural translation of k-variable formulas of the language of ρ_1 into formulas in the language of ρ_1°:

- If $\theta(x_1, ..., x_k)$ is a quantifier-free, then $\theta^\circ(v)$ is

$$\bigvee \{V_q(v) : \theta(x_{i_1}, ..., x_{i_m}) \in q(\overline{x})\}$$

- $(\theta_1 \vee \theta_2)^\circ = (\theta_1^\circ \vee \theta_2^\circ)$ and $(\neg\theta)^\circ = (\neg\,\theta^\circ)$
- $(\exists x_i\, \theta(\overline{x}))^\circ(v) = \exists u(E_{\{i\}}(u,v) \wedge \theta^\circ(u))$

The following, lemma 4.6, is folklore – proofs can be found in [19, 11, 1]. The succeeding assertion, proposition 4.2, is a straightforward consequence of lemma 4.6 and theorem 4.4. (Its proof is tedious and not at all enlightening, so we omit it.)

Lemma 4.6. *Let $\mathcal{A}$ be a ρ_1-structure. Then for any k-variable formula $\varphi(x_1,...,x_k)$ of the language of ρ_1, for any $\overline{a} \in A^k$*

$$\mathcal{A} \models \varphi(\overline{a}) \;\Leftrightarrow\; \mathcal{A}/P \models \varphi^\circ(\overline{a}/E^{\mathcal{A}})$$

Moreover, for any k-variable sentence ψ of ρ_1, $\mathcal{A} \models \psi$ if and only if $\mathcal{A}/P \models \psi^\circ$.

Proposition 4.7. *Let $\varphi(x_1,...,x_n)$ be a formula of ρ_1 such that $\varepsilon = (X \leftarrow \mathtt{invent}_k\{\overline{x}|\varphi\})$ is a basic expression of P. There is a sequence of basic expressions $\delta_1;...;\delta_m$ over ρ_1° such that for any ρ_1-structure $\mathcal{A}$,*

$$\mathtt{resp}(\varepsilon,\mathcal{A})/P = \mathtt{resp}(\delta_1;...;\delta_m,\mathcal{A}/P)$$

As a consequence of proposition , we find that there is program P° over ρ_1° of the form

$$P^\circ_{pre};(\mathbf{while}\,\varphi^\circ_{loop}\,\mathbf{do}\,(\mathbf{if}\,\psi_1^\circ\,\mathbf{then}\,R_1^\circ\,\mathbf{else}\,\mathrm{id});...;(\mathbf{if}\,\psi_m^\circ\,\mathbf{then}\,R_m^\circ\,\mathbf{else}\,\mathrm{id}))$$

which satisfies the following conditions for every ρ_1-structure $\mathcal{A}$:

1. $\mathtt{resp}(P^\circ_{pre},\mathcal{A}/P) = \mathtt{resp}(P_{pre},\mathcal{A})/P$ and $\mathtt{resp}(R_i^\circ,\mathcal{A}/P) = \mathtt{resp}(R_i,\mathcal{A})/P$ for each $i = 1,...,m$
2. $\mathcal{A} \models \varphi_{loop}$ if and only if $\mathcal{A}/P \models \varphi^\circ_{loop}$, and for each $i = 1,...,m$, $\mathcal{A} \models \psi_i$ if and only if $\mathcal{A}/P \models \psi_i^\circ$.

Moreover, if P is essentially inflationary, then P° is as well. One of our key theorems, then, is an immediate consequence of these observations:

Theorem 4.8. *Let K be a set of ρ_1-structures which is closed under isomorphism.*

1. *Suppose $f : \mathbb{N} \to \mathbb{N}$ is such that $\mathtt{cpx}(P,\mathcal{A}) \leq f(|A|)$ for all $\mathcal{A} \in K$, and suppose $g : \mathbb{N} \to \mathbb{N}$ is such that $g(n^k) \leq f(n)$ for all $n < \omega$. Then $\mathtt{cpx}(P^\circ,\mathcal{A}/P) \leq g(|A^k/E^{\mathcal{A}}|)$ for all $\mathcal{A} \in K$.*
2. *Suppose $g : \mathbb{N} \to \mathbb{N}$ is a function such that $\mathtt{cpx}(P^\circ,\mathcal{A}/P) \leq g(|A^k/E^{\mathcal{A}}|)$ for all $\mathcal{A} \in K$. Then $\mathtt{cpx}(P,\mathcal{A}) \leq g(|A|^k)$ for all $\mathcal{A} \in K$.*

In particular, P is polynomial-time if and only if P° is polynomial-time.

Given theorem 4.4, it is not difficult to verify that the transformation $\mathcal{A} \mapsto \mathcal{A}/P$ is computable in relational polynomial-time by an inflationary program Q_P. Now, we suppose that P *is* a polynomial-time `invent`-program (over K). Since $\mathcal{A}/P$ is a linearly ordered structure whenever $\mathcal{A} \in K$, we know that P° is captured over $K/P = \{\mathcal{A}/P : \mathcal{A} \in K\}$ by an essentially inflationary program, say $P^{\circ\circ}$ (this is by proposition 4.1 and lemma 4.3). Finally, there is a (constant-time) sequence of basic expressions, R, such that $Q_P; P^{\circ\circ}; R$ captures P. This line of argument is summarized in the following theorem:

Theorem 4.9. *Let K be a set of initial ρ_1-structures which is closed under isomorphism, and suppose Q is an HL-transformation of type $\rho_1 \to \rho_2$ such that $K \subseteq dom(Q)$. If Q is computable over K by a polynomial-time* `invent`*-program, then it is computable over K by an essentially inflationary polynomial-time* `invent`*-program.*

4.3 Efficiency and small algebraicity

Much of our later analysis of efficient model-building is insensitive to the distinction between genuinely polynomial-time `invent`-programs – which we've just learned are captured by essentially inflationary programs – and essentially inflationary programs that are not necessarily efficient. Dealing in program analysis alone, it does not seem likely that we will find a more refined characterization, and we, therefore, turn to model theory to provide that refinement.

For this section, we assume that K is an $\forall\exists$-axiomatized class of finite ρ-structures with amalgamation over sets. Suppose $\mathfrak{M}$ is a countable Fraïssé limit of K; in particular, $\mathfrak{M}$ is ultrahomogeneous, and $Th(\mathfrak{M})$ is $\aleph_0$-categorical and eliminates quantifiers. We further assume that P is a fully coherent solution of the model-building problem for K, and without loss of generality, we assume that $\mathtt{resp}(P, \mathcal{A}) \subseteq ||\mathfrak{M}||$ whenever $A \subset_{\mathbf{fin}} ||\mathfrak{M}||$, $\mathcal{A} = (A; \mathfrak{M})$. It follows that $acl^{\mathfrak{M}}(A) \subseteq \mathtt{resp}(P, \mathcal{A})$ in this case.

Lemma 4.10 (Polynomial-space)**.** *Suppose P is as stated above. If there are numbers δ, d such that $\mathtt{cpx}(P, (A; \mathfrak{M})) \leq \delta |A|^d$ for all sufficiently large $A \subset_{\mathbf{fin}} ||\mathfrak{M}||$, then are number δ_1, d_1 such that $\#||\mathtt{resp}(P, (A; \mathfrak{M}))|| \leq \delta_1 |A|^{d_1}$ for all large enough A.*

Proof. It is easy to see that if ε is a basic expression, then there is a number $r(\varepsilon)$ such that for any structure (of the appropriate signature) $\#||\mathtt{resp}(\varepsilon, \mathcal{A})|| \leq |A|^{r(\varepsilon)}$. Let r be the maximum $r(\varepsilon)$ for all basic expressions ε appearing in P. Then, clearly, $\#||\mathtt{resp}(P, \mathcal{A})|| \leq \delta |A|^{r^d}$, which suffices. $\square$

For any finite subset $C \subset_{\mathbf{fin}} ||\mathfrak{M}||$, let $S^{\mathrm{alg}}(C)$ denote the set of $p(x) \in S_1^{\mathrm{qf}}(C; \mathfrak{M})$ which are algebraic. Of course, if $a \in acl^{\mathfrak{M}}(C)$, then $qtp(a/C; \mathfrak{M}) \in S^{\mathrm{alg}}(C)$, so the following corollary is immediate from the polynomial-space constraint.

Corollary 4.11. *Suppose P is a polynomial-time fully coherent solution of the model-building problem for K. Then there are $d_{alg} < \omega$ and $\delta_{alg} \in \mathbf{R}$ such that $|S^{alg}(C)| \leq \delta_{alg}|C|^{d_{alg}}$ whenever $C \subset_{\boldsymbol{fin}} ||\mathfrak{M}||$ is sufficiently large.*

Recall that K is said to have *small algebraicity* if there is a number $d_K < \omega$ such that for all $C \subset_{\mathbf{fin}} ||\mathfrak{M}||$, if $p(x) \in S_1^{\mathrm{qf}}(C; \mathfrak{M})$ is algebraic, then $p(x)\restriction C_0$ is already algebraic for some $C_0 \in \binom{C}{d_K}$. Clearly, for any $C \subset_{\mathbf{fin}} ||\mathfrak{M}||$ and $0 < r < \omega$,

$$\left\{ p \in S_1^{\mathrm{qf}}(C) : \exists C_0 \in \binom{C}{r} . \, p\restriction C_0 \in S^{\mathrm{alg}}(C_0) \right\} \subseteq S^{\mathrm{alg}}(C)$$

Thus, if P satisfies a polynomial-time constraint, then for some $\varepsilon \in \mathbf{R}$, we have

$$\varepsilon |C|^{d_{\mathrm{alg}}} \leq \left| \left\{ p \in S_1^{\mathrm{qf}}(C) : \exists C_0 \in \binom{C}{r} . \, p\restriction C_0 \in S^{\mathrm{alg}}(C_0) \right\} \right| \leq |S^{\mathrm{alg}}(C)| \leq \delta_{\mathrm{alg}} |C|^{d_{\mathrm{alg}}}$$

whenever $C \subset_{\mathbf{fin}} ||\mathfrak{M}||$ is sufficiently large. That is, $|S^{\mathrm{alg}}(C)| \in \Theta(|C|^{d_{\mathrm{alg}}})$, and for some $c < \omega$ and all large enough C,

$$S^{\mathrm{alg}}(C) \subseteq \left\{ p \in S_1^{\mathrm{qf}}(C) : \exists C_0 \in \binom{C}{cd_{\mathrm{alg}}} . \, p\restriction C_0 \in S^{\mathrm{alg}}(C_0) \right\}$$

which proves:

Theorem 4.12. *Suppose $K = \boldsymbol{fin}[T^G]$, where T is a capped complete L^k-theory with infinitely many finite models up to isomorphism. If K is efficiently constructible – that is, K admits a* polynomial-time *fully coherent solution of its model-building problem – then K has small algebraicity.*

Chapter 5

þ-Independence and rosy Fraïssé classes

5.1 Definitions for þ-independence and rosiness

For easier reference in the sequel, we collect most of the fundamental definitions and notation together here. They are little changed from the classical case of [22]. (One difference is that there is nothing to distinguish types from formulas, as everything in sight is finite.) There are alternative but equivalent ways of making these definitions, but naturally, the ones we use suffice for our purposes.

Throughout, we take K to be a Fraïsse class – that is, K is an isomorphism-closed class of finite structures with infinitely many finite models up to isomorphism and which admits amalgamation over sets. Moreover, all types in consideration are assumed to be quantifier-free, though we will tend to indicate this explicitly with the notation *qtp*. Let $\mathcal{M}_0 \in K$ and $C \subseteq M_0$. Let $\varphi(x_1..., x_m, y_1, ..., y_n)$ be a partial (quantifier-free) type over $(C; \mathcal{M}_0)$, and let $q(\overline{y}) = qtp(\overline{b}_0/C; \mathcal{M}_0)$ for some $\overline{b}_0 \in M_0^n$. The first two definitions are used in all of our other definitions; the first of these, setwise unboundedness, obviates the need to employ imaginaries explicitly.

- The assertion, "$\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash q \,...\}$ is *setwise unbounded over* K_C,"[1] means that for every $t < \omega$, there is a model $\mathcal{M} \in K_C$ such that

  $$|\left\{\theta(M^m, \overline{b}) : \overline{b} \in M^n, \overline{b} \vDash q \,...\right\}| \geq t$$

  where (as is standard) $\varphi(M^n, \overline{b}) = \left\{\overline{a} \in M^n : \mathcal{M} \vDash \theta(\overline{a}, \overline{b})\right\}$.

- The assertion, "$\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash q \,...\}$ is *r-inconsistent over* K_C," means that for every $\mathcal{M} \in K_C$, if $\overline{b}_1, ..., \overline{b}_r \in q(M^n)$ are pairwise distinct, then

  $$\varphi(M^m, \overline{b}_1) \cap \cdots \cap \varphi(M^m, \overline{b}_r) = \emptyset$$

[1] The ellipsis "..." means that there may be additional constraints involved in the specification of the family of sets.

Now, we move to the fundamental definitions of the theory of þ-independence. Connecting these with our study of finite-variable theories – specifically, with the notion of game tableaux – we assume that every relation symbol R is of arity at least two and, of course, that the signature is finite.

- We say that $\varphi(\overline{x}, \overline{y})$ *divides strongly over C in q* if there is a positive integer r such that $\left\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash q\right\}$ is setwise unbounded and r-inconsistent over K_C.
- $\varphi(\overline{x}, \overline{b})$ divides strongly over C just in case it divides strongly over C in $qtp(\overline{b}/C; \mathcal{M}_0)$
- $\varphi(\overline{x}, \overline{b})$ *þ-divides over C* if there is a model $\mathcal{M}_1 \in K_C$ and a tuple $\overline{d} \in M_1^{<\omega}$ such that $\varphi(\overline{x}, \overline{b})$ divides strongly over $C\overline{d} = C \cup rng(\overline{d})$.
- We say that $\varphi(\overline{x}, \overline{b})$ *þ-forks over C* if there are $\mathcal{M} \in K_C$, formulas (i.e. partial types) $\theta_1(\overline{x}, \overline{y}_1)$, ..., $\theta_s(\overline{x}, \overline{y}_s)$ and $\overline{b}'_i \in M^{n_i}$, $1 \leq i \leq s < \omega$, such that

  $$K_{C\overline{b}\overline{b}_1 \ldots \overline{b}_s} \vDash \varphi(\overline{x}, \overline{b}) \to \bigvee_{i=1}^{s} \theta_i(\overline{x}, \overline{b}'_i)$$

  and each $\theta_i(\overline{x}, \overline{b}'_i)$ þ-divides over C.

Finally, we define the notion of þ-independence, $\mathop{\smile\!\!\!\!\!\!\mid}^{þ}$, as follows: Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$; then $\overline{a}$ *is þ-independent from B over C* – denoted $\overline{a} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_C B$ – just in case $qtp(\overline{a}/BC; \mathcal{M})$ does not þ-fork over C.

As is often the case in geometric model theory, there is a concept of an ordinal-valued of local rank associated with the notions þ-dividing and þ-forking. Again, let $\mathcal{M}_0 \in K$ and $C \subseteq M_0$, and let $\pi(x_1, ..., x_m)$ be a partial type over C. Let Φ be a a finite set of quantifier-free formulas without parameters of the form $\varphi(\overline{x}, \overline{y})$, and similarly, let Δ be a finite set of quantifier-free formulas of the form $\delta(\overline{y}, \overline{z})$.

- $þ_C(\pi, \Phi, \Delta, r) \geq 0$ just in case π is K_C-consistent.
- $þ_C(\pi, \Phi, \Delta, r) \geq \alpha + 1$ if there are formulas $\varphi(\overline{x}, \overline{y}) \in \Phi$ and $\delta(\overline{y}, \overline{z}) \in \Delta$, a model $\mathcal{M} \in K_C$ and a tuple $\overline{d} \in M^{\overline{z}}$ such that
  1. $\left\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash \delta(\overline{y}, \overline{d})\right\}$ is setwise unbounded and r-inconsistent over $K_{C\overline{d}}$
  2. $\left\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash \delta(\overline{y}, \overline{d}), þ_{C\overline{d}}(\pi \cup \{\varphi(\overline{x}, \overline{b})\}, \Phi, \Delta, r) \geq \alpha\right\}$ is setwise unbounded over $K_{C\overline{d}}$.
- $þ_C(\pi, \Phi, \Delta, r) = \infty$ if $þ_C(\pi, \Phi, \Delta, r) \geq \alpha$ for all $\alpha < \omega$.

  (For economy in the notation, we assert that $\alpha < \infty$ for all ordinals α.)

- We say that K *is rosy* if $\text{þ}_C(\pi, \Phi, \Delta, r) < \infty$ for all C, π, Φ, Δ as above.

The central result of this chapter is the demonstration that K is rosy if and only if $\mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}$ is a true independence relation in the sense specified in the Introduction (chapter 1). Before embarking on this development, we collect a few facts about the local

5.2 Basic properties of the local þ-rank

Lemma 5.1. *The local þ-rank has the following properties in all classes (with or without an amalgamation property):*

(Rank-monotonicity) Suppose $\Phi \subseteq \Phi'$, $\Delta \subseteq \Delta'$ and $\pi \subseteq \pi'$, where π' is a partial type over $(C; \mathcal{M}_0)$ and $\mathcal{M}_0 \in K$. Then

$$\text{þ}_C(\pi', \Phi', \Delta', r) \geq \text{þ}_C(\pi, \Phi, \Delta, r)$$

(Rank-transitivity) Suppose $\pi_0 \subseteq \pi_1 \subseteq \pi_2$, where π_2 is a partial type over $(C; \mathcal{M}_0)$ and $\mathcal{M}_0 \in K$. Assuming $\text{þ}_C(\pi_0, \Phi, \Delta, r) < \infty$, the following are equivalent:

1. $\text{þ}_C(\pi_0, \Phi, \Delta, r) = \text{þ}_C(\pi_2, \Phi, \Delta, r)$
2. $\text{þ}_C(\pi_0, \Phi, \Delta, r) = \text{þ}_C(\pi_1, \Phi, \Delta, r)$ *and* $\text{þ}_C(\pi_1, \Phi, \Delta, r) = \text{þ}_C(\pi_2, \Phi, \Delta, r)$

(Additivity) Let π_0, π_1 be quantifier-free formulas over $(C; \mathcal{M}_0)$ and $\mathcal{M}_0 \in K$. Then

$$\text{þ}_C(\pi_0 \vee \pi_1, \Phi, \Delta, r) = \max_{i=0,1} \text{þ}_C(\pi_i, \Phi, \Delta, r)$$

(Rank-extension) Let π be a partial type over $(C; \mathcal{M}_0)$ and $\mathcal{M}_0 \in K$. Then there is a complete (quantifier-free) type p over $(C; \mathcal{M}_0)$ extending π such that

$$\text{þ}_C(\pi, \Phi, \Delta, r) = \text{þ}_C(p, \Phi, \Delta, r)$$

proof of rank-monotonicity. Clearly, it will suffice to prove the following claim:

Claim. *If π, π', Φ, Φ' and Δ, Δ' are as in the statement of the lemma and $\alpha < \omega$, then $\text{þ}_C(\pi, \Phi, \Delta, r) \geq \alpha$ implies $\text{þ}_C(\pi', \Phi', \Delta', r) \geq \alpha$*

proof of claim. The proof is by induction on α. For $\alpha = 0$, it's enough to observe that since $\pi \supseteq \pi'$, if π is K_C-consistent, then π' must be also. Now, inductively assume that $\text{þ}_C(\pi, \Phi, \Delta, r) \geq \alpha_0$ implies $\text{þ}_C(\pi', \Phi', \Delta', r) \geq \alpha_0$ whenever $\alpha_0 \leq \alpha$.

Suppose that $\text{þ}_C(\pi, \Phi, \Delta, r) \geq \alpha + 1$ is witnessed by $\varphi(\overline{x}, \overline{y}) \in \Phi$, $\delta(\overline{y}, \overline{z}) \in \Delta$, $\mathcal{M} \in K_C$ and $\overline{d} \in M^{\overline{z}}$. Let $t < \omega$ be given. Then, there is a model $\mathcal{M}_1 \in K_{C\overline{d}}$ such that

$$\left\{\varphi(\overline{x}, \overline{b}) : \mathcal{M}_1 \models \delta(\overline{b}, \overline{d}),\ \text{þ}_{C\overline{d}}(\pi \cup \{\varphi(\overline{x}, \overline{b})\}, \Phi, \Delta, r) \geq \alpha\right\}$$

has cardinality at least t (as a family of subsets of $M_1^{\overline{x}}$). By the inductive hypothesis, it follows that

$$\left\{\varphi(\overline{x},\overline{b}) : \mathcal{M}_1 \vDash \delta(\overline{b},\overline{d}),\ þ_{C\overline{d}}(\pi' \cup \{\varphi(\overline{x},\overline{b})\}, \Phi', \Delta', r) \geq \alpha\right\}$$

Of course, $\{\varphi(\overline{x},\overline{b}) : \mathcal{M}_1 \vDash \delta(\overline{b},\overline{d})\}$ remains r-inconsistent. As t was arbitrary, φ, δ, $\mathcal{M}_1$ and $\overline{d}$ witness the fact that $þ_C(\pi', \Phi', \Delta', r) \geq \alpha + 1$, as required. □

□

proof of rank-transitivity. By rank-monotonicity, we find that

$$þ_C(\pi_2, \Phi, \Delta, r) \leq þ_C(\pi_1, \Phi, \Delta, r) \leq þ_C(\pi_0, \Phi, \Delta, r) < \infty$$

and rank-transitivity follows immediately from the inequality. □

proof of additivity. In view of rank-monotonicity, it suffices to prove:

Claim. *For every $\alpha < \omega$,*

$$þ_C(\pi_0 \vee \pi_1, \Phi, \Delta, r) \geq \alpha \ \Rightarrow\ \max_{i=0,1} þ_C(\pi_i, \Phi, \Delta, r) \geq \alpha$$

Proof.

Sub-claim. *Suppose $þ_C(\pi_0 \vee \pi_1, \Phi, \Delta, r) \geq \alpha + 1$, witnessed by $\varphi(\overline{x},\overline{y}) \in \Phi$, $\delta(\overline{y},\overline{z}) \in \Delta$, $\mathcal{N} \in K_C$ and $\overline{d} \in N^{\overline{z}}$. Then, for every $m < \omega$, there are $\mathcal{N}_1 \in K_{C\overline{d}}$ and an $i < 2$ such that*

$$\left\{\varphi(\overline{x},\overline{b}) : \mathcal{N}_1 \vDash \delta(\overline{b},\overline{d}),\ þ_{C\overline{d}}(\pi_i \cup \{\varphi(\overline{x},\overline{b})\}, \Phi, \Delta, r) \geq \alpha\right\}$$

has cardinality at least m (as a family of subsets of $N_1^{\overline{x}}$).

proof of subclaim. Let $m < \omega$ be given. Choose any $n \geq 2m+1$, and let $\mathcal{N}_m \in K_{C\overline{d}}$ be such that

$$H_m \left\{\varphi(\overline{x},\overline{b}) : \mathcal{N}_m \vDash \delta(\overline{b},\overline{d}),\ þ_{C\overline{d}}((\pi_0 \vee \pi_1) \cup \{\varphi(\overline{x},\overline{b})\}, \Phi, \Delta, r) \geq \alpha\right\}$$

has cardinality at least n as a family of subsets of $N_m^{\overline{x}}$. Note that for every $\overline{b} \in N_m^{\overline{y}}$, if $\varphi(\overline{x},\overline{b}) \in H_m$, then at least one of $\pi_0 \cup \{\varphi(\overline{x},\overline{b})\}$ or $\pi_1 \cup \{\varphi(\overline{x},\overline{b})\}$ is $K_{C\overline{db}}$-consistent. Define $h_m : H_m \to 2$ by

$$h_m(\varphi(\overline{x},\overline{b})) = \min\left\{i < 2 : \pi_0 \cup \{\varphi(\overline{x},\overline{b})\} \text{ is } K_{C\overline{db}}\text{-consistent}\right\}$$

By our choice of n, either $|h_m^{-1}(0)| \geq m$ or $|h_m^{-1}(1)| \geq m$, and obviously this proves the sub-claim. □

Assume $þ_C(\pi_0 \vee \pi_1, \Phi, \Delta, r) \geq \alpha + 1$, witnessed by $\varphi(\overline{x}, \overline{y}) \in \Phi$, $\delta(\overline{y}, \overline{z}) \in \Delta$, $\mathcal{N} \in K_C$ and $\overline{d} \in N^{\overline{z}}$. Define $f : \omega \to 2$ by

$$f(m) = \min \left\{ i < 2 : |h_m^{-1}(i)| \geq m \right\}$$

where $\mathcal{N}_m$, H_m and h_m are those chosen in the proof of the sub-claim. At least one of $f^{-1}(0)$ or $f^{-1}(1)$ is infinite, and in case $f^{-1}(i)$ is infinite, φ, δ, $\mathcal{N}$, $\overline{d}$ witness the fact that $þ_C(\pi, \Phi, \Delta, r) \geq \alpha + 1$. □

□

proof of rank-extension. Assume $\pi = \pi(x_1, ..., x_n)$. Since ρ and C are finite, we may enumerate $p_0, ..., p_{s-1}$ all of the complete extensions of π to C. Clearly, $K_C \vDash \pi \leftrightarrow \bigvee_{i<s} p_i$, so by rank-additivity, we have,

$$þ_C(\pi, \Phi, \Delta, r) = \max_{i<s} þ_C(p_i, \Phi, \Delta, r)$$

and any type p_i realizing maximum is a witness to rank-extension. □

Proposition 5.2 (Invariance of þ-rank)**.** *Assume K has amalgamation over sets. Let $\mathcal{M} \in K$ and $C, C' \subseteq M$, and let π be a partial type over $C \cap C'$. Then*

$$þ_C(\pi, \Phi, \Delta, r) = þ_{C'}(\pi, \Phi, \Delta, r)$$

Proof. As usual, the proposition reduces to the following claim.

Claim. *Let $\alpha < \omega$. For any Φ and Δ, and for any $\mathcal{M} \in K$, any $C, C' \subseteq M$, and be any partial type π over $C \cap C'$, $þ_C(\pi, \Phi, \Delta, r) \geq \alpha$ implies $þ_{C'}(\pi, \Phi, \Delta, r) \geq \alpha$.*

proof of claim. Suppose that $þ_C(\pi, \Phi, \Delta, r) \geq \alpha + 1$, witnessed by $\varphi(\overline{x}, \overline{y}) \in \Phi$, $\delta(\overline{y}, \overline{z}) \in \Delta$, $\mathcal{M}_1 \in K_C$ and $\overline{d} \in M_1^{\overline{z}}$. Let $\mathcal{M}_2 \in K_{C'}$, and let $\mathcal{N}$ be an amalgam of $\mathcal{M}_1$ and $\mathcal{M}_2$ over $C \cap C'$. Then it's easy to see that (up to the embeddings of the amalgamation) $\varphi(\overline{x}, \overline{y}) \in \Phi$, $\delta(\overline{y}, \overline{z}) \in \Delta$, $\mathcal{M}_1 \in K_C$ and $\overline{d} \in N^{\overline{z}}$. □

□

5.3 Basic properties of þ-independence

In this section, we collect the basic properties of $\mathop{\smile\!\!\!\!\!\!|}^{þ}$ (proposition 5.3) that hold in any class (with AP/sets), rosy or not. That is, we show that þ-independence is always a notion of independence. In the next section, we will find that $\mathop{\smile\!\!\!\!\!\!|}^{þ}$ is a true notion of independence in a class K precisely when K is rosy.

Proposition 5.3. *The relation $\mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}$ has the following properties in any class K with amalgamation over sets, rosy or not.*

1. Invariance*: Suppose $\mathcal{M}, \mathcal{M}_1 \in K$, $B, C \subseteq M$, $B_1, C_1 \subseteq M_1$ and $\overline{a} \in M^{<\omega}$, $\overline{a}_1 \in M_1^{<\omega}$, and suppose $qtp(\overline{a}BC; \mathcal{M}) = qtp(\overline{a}_1 B_1 C_1; \mathcal{M}_1)$. If $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_B C$, then $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_{C_1} B_1$.*
2. Extension*: Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C B$. Let $\mathcal{M}' \in K_{BC}$, and suppose $BC \subseteq D \subseteq M'$.*

   Then there are $\mathcal{M}_1 \in K_D$ and $\overline{a}_1 \in M_1^{<\omega}$ such that $qtp(\overline{a}_1/BC; \mathcal{M}_1) = qtp(\overline{a}/BC; \mathcal{M})$ and $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C D$.
3. Base-extension *Let $\mathcal{M}_0 \in K$, $B \subseteq M_0$ and $\overline{a} \in M_0^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C B$. Let $\mathcal{M} \in K_{BC}$, and suppose $\overline{c} \in M^{<\omega}$.*

   Then there are $\mathcal{M}_1 \in K_{\overline{a}BC}$ and $\overline{c}' \in M_1^{<\omega}$ such that $qtp(\overline{c}'/\overline{a}BC; \mathcal{M}) = qtp(\overline{c}/\overline{a}BC; \mathcal{M})$ and $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_{C\overline{c}'} B$
4. Monotonicity*: Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C B$. If $B_0 \subseteq B$, then $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C B_0$.*
5. Base-monotonicity*: Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C B$. If $B_0 \subseteq B$, then $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_{B_0 C} B$.*
6. Partial right-transitivity *Let $\mathcal{M} \in K$, $B_1, B_2, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C B_1 B_2$. Then $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C B_1$ and $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_{B_1 C} B_2$.*
7. Preservation of algebraic dependence I*: Suppose $\mathcal{M} \in K$, $B \subseteq M$ and $\overline{a} \in M^{<\omega}$. If $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_B \overline{a}$, then $\overline{a} \in acl(B; \mathcal{M})$.*
8. Preservation of algebraic dependence II*: Suppose $\mathcal{M} \in K$, $C \subseteq M$ and $\overline{a}, \overline{b} \in M^{<\omega}$. If $\overline{b} \in acl(C\overline{a}) \setminus acl(C)$, then $\overline{a} \not\mathop{\smile\!\!\!\!\!\!\downarrow}^{þ}_C \overline{b}$.*

proof of extension. For $\mathcal{M} \in K$ and $B_0 \subseteq B \subseteq M$, define

$$D^{þ}_{\overline{x}}(B/B_0) = \left\{\varphi(\overline{x}, \overline{b}) : \varphi(\overline{x}, \overline{y}) \text{ þ-divides over } B_0\right\}$$
$$nD^{þ}_{\overline{x}}(B/B_0) = \left\{\neg\varphi(\overline{x}, \overline{b}) : \varphi(\overline{x}, \overline{b}) \in D^{þ}_{\overline{x}}(B/B_0)\right\}$$

Claim. *$\mathcal{M} \in K$ and $B_0 \subseteq B \subseteq M$, and let $\pi(\overline{x})$ be a partial type over C. Then, the partial type $\pi(\overline{x}) \cup \mathrm{nD}^{þ}_{\overline{x}}(B/B_0)$ is K-realizable if and only if $\pi(\overline{x})$ does not þ-fork over B_0.*

Proof. ($\Leftarrow$) For the contrapositive, suppose $\pi(\overline{x}) \cup nD^{þ}_{\overline{x}}(B/B_0)$ is not K-realizable. Let $p_1(\overline{x}), ..., p_s(\overline{x})$ be an enumeration of all of the complete extensions $\pi(\overline{x})$ to B. By hypothesis, $p_i(\overline{x}) \cap D^{þ}_{\overline{x}}(B/B_0)$ is non-empty for each $i \in [s]$, and it follows immediately that $\pi(\overline{x})$ þ-forks over B_0.
($\Rightarrow$) Immediate from AP/sets and the definition of þ-forking. $\square$

Now, let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B$. Let $\mathcal{M}' \in K_{BC}$, and suppose $BC \subseteq D \subseteq M'$. Let $p(\overline{x}) = qtp(\overline{a}/BC; \mathcal{M})$, which does not þ-fork over BC. By the claim, there is a complete extension $p'(\overline{x})$ of $p(\overline{x}) \cup nD^{\text{þ}}_{\overline{x}}(D/C)$ to D, and by definition of $nD^{\text{þ}}_{\overline{x}}(D/C)$, $p'(\overline{x})$ does not þ-fork over C, as desired. □

proof of monotonicity. Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B$ and $B_0 \subseteq B$. Let $\overline{b}$ enumerate $B \setminus B_0$, and let $p(\overline{x}, \overline{y}) = qtp(\overline{a}, \overline{b}/B_0C; \mathcal{M})$ and $p_0(\overline{x}) = qtp(\overline{a}/B_0C; \mathcal{M})$. Towards a contradiction, suppose that $p_0(\overline{x})$ þ-forks over C. Then, there are $\overline{c}_1, ..., \overline{c}_s$ and formulas $\varphi_1(\overline{x}, \overline{y}_1)$, ..., $\varphi_s(\overline{x}, \overline{y}_s)$ such that

$$K_{C\overline{c}_1 \ldots \overline{c}_s} \vDash p_0(\overline{x}) \to \bigvee_{i=1}^{s} \varphi_i(\overline{x}, \overline{c}_i)$$

and each $\varphi_i(\overline{x}, \overline{c}_i)$ þ-divides over C. Clearly,

$$K_{C\overline{b}\overline{c}_1 \ldots \overline{c}_s} \vDash p(\overline{x}, \overline{b}) \to \bigvee_{i=1}^{s} \varphi_i(\overline{x}, \overline{c}_i)$$

so $\varphi_i(\overline{x}, \overline{c}_1)$, ..., $\varphi_s(\overline{x}, \overline{c}_s)$ witness the fact that $p(\overline{x}, \overline{b})$ þ-forks over C, contradicting the assumption that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B$. □

proof of base-monotonicity. Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B$ and $B_0 \subseteq B$. Let $p(\overline{x}) = qtp(\overline{a}/BC; \mathcal{M})$. Towards a contradiction, suppose $\overline{a} \not\mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{B_0C} B$ – that is, let $\varphi_1(\overline{x}, \overline{y}_1)$, ..., $\varphi_s(\overline{x}, \overline{y}_s)$ be formulas and $\overline{c}_1, ..., \overline{c}_s$ tuples such that

$$K_{B_0C\overline{c}_1 \ldots \overline{c}_s} \vDash p(\overline{x}) \to \bigvee_{i=1}^{s} \varphi_i(\overline{x}, \overline{c}_i)$$

and each $\varphi_i(\overline{x}, \overline{c}_i)$ þ-divides over B_0C. By definition, then, $\varphi_i(\overline{x}, \overline{c}_i)$ þ-divides over C as well, so these formulas witness $\overline{a} \not\mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B$, a contradiction. □

proof of partial right-transitivity. Let $\mathcal{M} \in K$, $B_1, B_2, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B_1B_2$. By Monotonicity, we know that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B_1$, so it remains only to show that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{B_1C} B_2$. Again by Monotonicity, it suffices to show that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{B_1C} B_1B_2$, and this follows from $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B_1B_2$ by Base-monotonicity, which completes the proof of the claim. □

proof of preservation of algebraic dependence I. Suppose $\mathcal{M} \in K$, $B \subseteq M$ and $\overline{a} \in M^l$. Suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{B} \overline{a}$ and $\overline{a}$ is not algebraic. For a contradiction, we will show that $qtp(\overline{a}/\overline{a}B; \mathcal{M})$ divides strongly (hence þ-forks) over B. Naturally, we let $\varphi(x_1, ..., x_l, y_1, ..., y_l)$ be the formula $\bigwedge_{i=1}^{l} x_i = y_i$, and let $q(\overline{x}) = qtp(\overline{a}/B; \mathcal{M})$. Then, as $q(\overline{x})$ is not algebraic,

$$\{\varphi(\overline{x}, \overline{a}') : \overline{a}' \vDash q\}$$

is setwise unbounded and 2-inconsistent over K_B, as desired. □

proof of preservation of algebraic dependence II. We name the complete types $p(\overline{x}, \overline{y}) = qtp(\overline{a}, \overline{b}/C; \mathcal{M})$, $p_0(\overline{x}) = qtp(\overline{a}/C; \mathcal{M})$ and $p_1(\overline{y}) = qtp(\overline{b}/C; \mathcal{M})$. Since $acl(C\overline{a}) \neq acl(C)$, we know that p_0 is non-algebraic, and p_1 is non-algebraic by hypothesis. But, as $p(\overline{a}, \overline{y})$ is algebraic, let $0 < r < \omega$ be such that

$$|\left\{p(\overline{a}, \overline{b}') : \overline{b}' \in N^{\overline{y}}\right\}| < r$$

as $\mathcal{N}$ ranges over $K_{C\overline{a}}$. It is easy enough to see that if $\mathcal{M}_1 \in K_C$ and $\overline{a}_1 \in M_1^{\overline{x}}$ is such that $\overline{a}_1 \models p_0$, then

$$|\left\{p(\overline{a}_1, \overline{b}') : \overline{b}' \in N^{\overline{y}}\right\}| < r$$

as $\mathcal{N}$ ranges over $K_{C\overline{a}_1}$. As p_1 is non-algebraic, it follows that

$$\left\{p(\overline{x}, \overline{b}') : \overline{b}' \models p_1\right\}$$

is both setwise unbounded and r-inconsistent over K_C; that is, $p(\overline{a}, \overline{y})$ þ-divides – hence, þ-forks – over C, as required. □

proof of base-extension. Let $\mathcal{M}_0 \in K$, $B \subseteq M_0$ and $\overline{a} \in M_0^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} B$. Let $\mathcal{M} \in K_{BC}$, and suppose $\overline{c} \in M^{<\omega}$. . Let $p(\overline{x}, \overline{y}) = qtp(\overline{ac}/BC; \mathcal{M})$. By Extension, we can choose an $\mathcal{M}_1 \in K_{BC\overline{c}}$ and an $\overline{a}_1 \in M_1^{<\omega}$ such that $\mathcal{M}_1 \models p(\overline{a}_1, \overline{c})$ and $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} B\overline{c}$, and by Partial right-transitivity, it follows that $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{c}} B$. Now, we may choose $\mathcal{M}_2 \in K_{BC\overline{a}}$ and $\overline{c}' \in M_2^{|\overline{c}|}$ such that $\mathcal{M}_2 \models p(\overline{a}, \overline{c}')$ By Invariance, then, we have $\overline{a} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{c}'} B$, as desired. □

Lemma 5.4 (Partial left-transitivity). *Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a}_1, \overline{a}_2 \in M^{<\omega}$, and suppose $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} B$ and $\overline{a}_2 \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{\overline{a}_1 C} B$. Then $\overline{a}_1\overline{a}_2 \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} B$.*

Proof. The proof is identical to that of the result of the same name in [22], so we omit the demonstration here. □

5.4 þ-Independence in rosy classes

In this section, we will show that þ-independence in a rosy class is a true independence relation. A key step in this development is the connection between the global property of þ-forking with the local property described by the local þ-rank. Thus, one main theorem of this section is the following:

Theorem 5.5. *Assume K is rosy. Let $\mathcal{M} \in K$ and $B, C \subseteq M$, and let $p(\overline{x})$ be a complete type over $(BC; \mathcal{M})$. Then, the following are equivalent:*

1. *$þ(p, \varphi, \delta, r) = þ(p{\restriction}C, \varphi, \delta, r)$ for all quantifier-free formulas $\varphi(\overline{x}, \overline{y})$ and $\delta(\overline{y}, \overline{z})$;*

2. p does not þ-fork over C.

We break the proof of theorem 5.5 into two propositions, 5.6 and 5.10 which are treated in this section.

Proposition 5.6. *Assume K is rosy. Let $\mathcal{M} \in K$ and $B, C \subseteq M$, and let $p(\overline{x})$ be a complete quantifier-free type over $(BC; \mathcal{M})$. Suppose $p(\overline{x})$ þ-forks over C. Then $þ(p, \varphi, \Delta, r) < þ(p{\restriction}C, \varphi, \Delta, r)$ for some $1 < r < \omega$, $\varphi(\overline{x}, \overline{y})$ and Δ.*

Proof. Assuming $p(\overline{x})$ þ-forks over C, we find $\overline{b}_1, ..., \overline{b}_s \in M^{<\omega}$ and formulas $\varphi_1(\overline{x}, \overline{y}_1)$,..., $\varphi_s(\overline{x}, \overline{y}_s)$ such that

$$K_{C\bar{b}_1 \dots \bar{b}_s} \vDash p(\overline{x}) \rightarrow \bigvee_{i=1}^{s} \varphi_i(\overline{x}, \overline{b}_s)$$

and for each $i \in [s]$, $\varphi_i(\overline{x}, \overline{b}_i)$ þ-divides over $(C; \mathcal{M})$. Without loss of generality, we may assume that $\overline{b}_1, ..., \overline{b}_s \in (BC)^{<\omega}$. Thus, for each $i \in [s]$, we may further assume that there is a tuple $\overline{d}_i \in M^{<\omega}$ such that $\varphi_i(\overline{x}, \overline{b}_i)$ divides strongly over $C\overline{d}_i$, and let $1 < r_i < \omega$ such that

$$\left\{\varphi_i(\overline{x}, \overline{b}) : \overline{b} \vDash qtp(\overline{b}_i / C\overline{d}_i; \mathcal{M})\right\}$$

is setwise unbounded and r_i-inconsistent over $K_{C\bar{d}_i}$. For convenience, let $\overline{c}$ be an enumeration of C. Then if $\delta_i(\overline{y}_i, \overline{v}, \overline{z}_i) = qtp(\overline{b}_i \overline{c} \overline{d}_i; \mathcal{M})$, we may also say that $\varphi_i(\overline{x}, \overline{y}_i)$ divides strongly in $\delta_i(\overline{y}_i, \overline{c}, \overline{d}_i)$ over $(C; \mathcal{M})$. For convenience in the exposition we will assume that $\overline{b}_1, ..., \overline{b}_s$ are all of the same arity and $\overline{d}_1, ..., \overline{d}_s$ are all of the same arity (possibly with repetitions of coordinates); in particular, we assume that $\overline{z}_1 = \cdots = \overline{z}_s$. We make several additional definitions

- $r = \max_{i \in [s]} r_i$
- $\overline{w}$ is a tuple of fresh variables of the same length as the $\overline{y}_i$'s, and $\overline{y}' = \overline{y}_1 \frown \cdots \frown \overline{y}_s \frown \overline{w}$; similarly, for each $i \in [s]$, $\overline{b}^i = \overline{b}_1 \frown \cdots \frown \overline{b}_s \frown \overline{b}_i$
- $\psi(\overline{x}, \overline{y}') = \bigvee_{i=1}^{s} (\overline{w} = \overline{y}_i \wedge \varphi_i(\overline{x}, \overline{y}_i))$
- $\delta_i'(\overline{y}', \overline{v}, \overline{z}) = \overline{w} = \overline{y}_i \wedge \delta_i(\overline{w}, \overline{v}, \overline{z})$
- $\Delta' = \{\delta_i'(\overline{y}', \overline{z}) : i \in [s]\}$

Clearly, if $\overline{a}' = \overline{a}_1 \frown \cdots \frown \overline{a}_s \frown \overline{a}_i$, then $\psi(\overline{x}, \overline{a}') \leftrightarrow \varphi_i(\overline{x}, \overline{a}_i)$ is valid over K. It only remains to prove,

Claim. $þ(p, \psi, \Delta', r) < þ(p{\restriction}C, \psi, \Delta', r)$

proof of claim. By construction, we know that

$$K_{C\bar{b}_1...\bar{b}_s} \vDash p(\overline{x}) \to \bigvee_{i=1}^{s} \psi(\overline{x}, \overline{b}^i)$$

and further, for each $i \in [s]$, $\left\{\psi(\overline{x}, \overline{b}') : \overline{b} \vDash \delta_i'(\overline{y}', \overline{c}, \overline{d}_i)\right\}$ is both setwise unbounded and r-inconsistent over $K_{C\bar{b}_1...\bar{b}_s}$. By rank-extension, there is a complete type $q(\overline{x})$ over the enlarged domain $(BC\overline{b}_1...\overline{b}_s; \mathcal{M})$ extending p such that $þ(q, \psi, \Delta', r) = þ(p, \psi, \Delta', r)$. Since

$$K_{C\bar{b}_1...\bar{b}_s} \vDash p(\overline{x}) \to \bigvee_{i=1}^{s} \psi(\overline{x}, \overline{b}^i)$$

we may assume that $\psi(\overline{x}, \overline{b}^1) \in q$, and clearly, $q{\restriction}C = p{\restriction}C$ by definition. Now,

$$\left\{\psi(\overline{x}, \overline{b}) : \overline{b} \vDash \delta_1'(\overline{y}', \overline{c}, \overline{d}_1)\right\}$$

is setwise unbounded and r-inconsistent over $K_{C\overline{d}_1}$ just because $\varphi_1(\overline{x}, \overline{d}_1)$ divides strongly over $(C; \mathcal{M})$ in $\delta_1(\overline{y}_1, \overline{c}, \overline{d}_1)$. By the definition of þ-rank, therefore, we have

$$\begin{aligned} þ(p{\restriction}C, \psi, \Delta', r) &= þ(q{\restriction}C, \psi, \Delta', r) \\ &\geq þ(q, \psi, \Delta', r) + 1 \\ &= þ(p, \psi, \Delta', r) + 1 \\ &> þ(p, \psi, \Delta', r) \end{aligned}$$

which proves the claim. □

□

Having proved proposition 5.6, we obtain a second proof of the Extension property of $\mathop{\smile\!\!\!\!\!\!|}^{þ}$ in rosy classes, which the author finds somewhat more natural to this scenario where everything in sight is finite. We also easily derive the Existence property of $\mathop{\smile\!\!\!\!\!\!|}^{þ}$ from proposition 5.6.

Observation (Extension in rosy classes)**.** *Let $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{þ}_C B$. Let $\mathcal{M}' \in K_{BC}$, and suppose $BC \subseteq D \subseteq M'$.*

Proof. We first choose a function $f_{\mathrm{rk}} : \omega \times \omega \to \omega$ such that for any $n < \omega$, $\mathcal{N} \in K$, $B_0 \subseteq B \subseteq N$, and any complete n-type $p(\overline{x})$ over $(B; \mathcal{N})$, the following are equivalent:

1. $þ(p, \Phi, \Delta, r) < þ(p{\restriction}B_0, \Phi, \Delta, r)$ for some Φ, Δ and r.

2. With $m = f_{\mathrm{rk}}(n, |B_0|)$, $þ(p, F_m, F_m, m) < þ(p{\restriction}B_0, F_m, F_m, m)$ where F_m is the set of $\leq m$-ary quantifier-free formulas, counted up to logical equivalence.

Such a function exists just due to the fact that, for all $n, t < \omega$, the set

$$\left\{(q, (A; \mathcal{N})) : \mathcal{N} \in K,\ A \in \binom{N}{t}, q(x_1, ..., x_n) \in S_n^{\mathrm{qf}}(A; \mathcal{N})\right\}$$

is finite up to isomorphism of the base sets $(A; \mathcal{N})$. We may also assume that f_{rk} is monotone in each coordinate – that is, if $n \leq n_1$ and $t \leq t_1$, then $f_{\mathrm{rk}}(n, t) \leq f_{\mathrm{rk}}(n_1, t_1)$.

Now, let $\mathcal{M} \in K$, $B, C, D \subseteq M$ such that $BC \subseteq D$, and $\overline{a} \in M^n$, and assume that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_C B$. Let $p_1(\overline{x}) = qtp(\overline{a}/BC; \mathcal{M})$ and $p_0 = p_1 \restriction C$, and let $m = f_{\mathrm{rk}}(n, |BC|)$. By Rank-extension, there is an extension $p_2(\overline{x})$ of $p_1(\overline{x})$ to $(D; \mathcal{M})$ such that

$$\text{þ}(p_2, F_m, F_m, m) = \text{þ}(p_1, F_m, F_m, m)$$

where $m = f_{\mathrm{rk}}(n, |BC|)$. To prove the claim, it suffices to show that $p_2(\overline{x})$ does not þ-fork over $(C; \mathcal{M})$, and by proposition 5.6, p_2 does not þ-fork over BC. Again, by proposition 5.6 and the monotonicity of f_{rk}, in order to show that p_2 does not þ-fork over C, it suffices to show that

$$\text{þ}(p_2, F_m, F_m, m) = \text{þ}(p_0, F_m, F_m, m)$$

Thus – towards a contradiction – we suppose that

$$\alpha = \text{þ}(p_1, F_m, F_m, m) = \text{þ}(p_2, F_m, F_m, m) < \text{þ}(p_0, F_m, F_m, m)$$

There are, then, $\varphi(\overline{x}, \overline{y}), \delta(\overline{y}, \overline{z}) \in F_m$ and (w.l.o.g.) $\overline{e} \in M^{\overline{z}}$ such that

$$\left\{\varphi(\overline{x}, \overline{b}') : \overline{b}' \vDash \delta(\overline{y}, \overline{e})\right\}$$

is m-inconsistent over $K_{BC\overline{e}}$ and

$$\left\{\varphi(\overline{x}, \overline{b}') : \overline{b}' \vDash \delta(\overline{y}, \overline{e}),\ \text{þ}(p_0 \cup \{\varphi(\overline{x}, \overline{b}')\}, F_m, F_m, m) \geq \alpha\right\}$$

is setwise unbounded. By AP/sets, some $\overline{b}' \vDash \delta(\overline{y}, \overline{e})$ is algebraic over $dom(p_1) = BC$, and it follows that p_1 þ-forks (even þ-divides) over C, which is a contradiction to the assumption that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_C B$. □

Proposition 5.7 (Existence property of $\mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}$ in rosy classes)**.** *Suppose $\mathcal{M} \in K$, $B \subseteq M$ and $\overline{a} \in M^{<\omega}$. Then $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_B B$.*

Proof. Immediate from proposition 5.6. □

The proof of the next theorem is mostly the same as the proof of the analogous theorem in [22] – the only novel aspect being the use of coherent sequences. We note that because we make recourse to the machinery of coherent sequences, the assumption of amalgamation over sets is not actually necessary in the argument, except insofar as it is used in proving the Extension property.

Theorem 5.8 (Symmetry). *Assume K is rosy. Let $M_0 \in K$, $C \subseteq M_0$, $\overline{a}, \overline{b} \in M_0^{\leq k}$. Then $\overline{a} \mathop{\smile\hskip -0.9em ^| \ }^{þ}_{C} \overline{b}$ implies $\overline{b} \mathop{\smile\hskip -0.9em ^| \ }^{þ}_{C} \overline{a}$.*

Proof. Suppose $M_0 \in K$, $C \subseteq M_0$ and $\overline{a} \in M_0^{\leq k}$, $\overline{b} \in M_0^{<\omega}$, and assume $\overline{a} \mathop{\smile\hskip -0.9em ^| \ }^{\text{þ}}_{C} \overline{b}$ but $\overline{b} \mathop{\not\smile\hskip -0.9em ^| \ }^{\text{þ}}_{C} \overline{a}$. We will derive a contradiction to the assumption that K is rosy. Let $p(\overline{x}, \overline{y}) = qtp(\overline{a}, \overline{b}/C; \mathcal{M}_0)$ – so, $p(\overline{a}, \overline{y})$ þ-forks over C. By definition, we can assume that there are formulas $\varphi_1(\overline{y}, \overline{v}_1)$, ..., $\varphi_s(\overline{y}, \overline{v}_s)$, tuples $\overline{e}_1, ..., \overline{e}_s$ and $\overline{d}_1, ..., \overline{d}_s$ in M_0 such that

$$K_{C\overline{a}\overline{e}_1...\overline{e}_s} \models p(\overline{a}, \overline{y}) \to \bigvee_{i=1}^{s} \varphi_i(\overline{y}, \overline{e}_i)$$

and for each $i \in [s]$, $\varphi_i(\overline{y}, \overline{e}_i)$ divides strongly over $(C\overline{d}_i; \mathcal{M}_0)$.

Our first goal is to define a coherent sequence $(\mathcal{N}_n, (\overline{a}_n, \overline{e}^n, \overline{d}^n))_{n<\omega}$ in the type

$$qtp(\overline{a}, \overline{e}^0, \overline{d}^0/C\overline{b}; \mathcal{M}_0)$$

(where $\overline{e}^0 = \overline{e}_1{}^\frown \cdots {}^\frown \overline{e}_s$ and $\overline{d}^0 = \overline{d}_1{}^\frown \cdots {}^\frown \overline{d}_s$) such that for all $n < \omega$, $\overline{a}_n \mathop{\smile\hskip -0.9em ^| \ }^{\text{þ}}_{C} \overline{b}\overline{a}_{<n}\overline{e}^{<n}\overline{d}^{<n}$.

Construction:

- Naturally, we set $\overline{a}_0 = \overline{a}$, $\overline{e}^0 = \overline{e}_1{}^\frown \cdots {}^\frown \overline{e}_s$ and $\overline{d}^0 = \overline{d}_1{}^\frown \cdots {}^\frown \overline{d}_s$. Let $p_0(\overline{x}) = qtp(\overline{a}/C; \mathcal{M}_0)$ and $q_0(\overline{y}, \overline{z}) = qtp(\overline{e}^0\overline{d}^0/C\overline{b}; \mathcal{M}_0)$.

- Suppose we have chosen $\overline{a}_i, \overline{e}^i, \overline{d}^i$ and $\mathcal{M}_i$ for all $i \leq n$, and the associated types are $p_i(\overline{x}) = qtp(\overline{a}_i/C\overline{a}_{<i}\overline{e}^{<i}\overline{d}^{<i}; \mathcal{M}_i)$ and $q_i = (\overline{e}^i, \overline{d}^i/C\overline{b}\overline{e}^{<i}\overline{d}^{<i}; \mathcal{M}_i)$. We also assume that p_i does not þ-fork over C.

  1. Applying Extension, choose $p_{n+1}(\overline{x})$ over $(C\overline{a}_{\leq n}\overline{e}^{\leq n}\overline{d}^{\leq n}; \mathcal{M}_n)$ extending p_n which does not þ-fork over C, and choose $\mathcal{M}'_{n+1} \in K_{C\overline{a}_{\leq n}\overline{e}^{\leq n}\overline{d}^{\leq n}}$ and $\overline{a}_{n+1}$ in $\mathcal{M}'_{n+1}$ realizing p_{n+1}.
  2. Applying Extension again, choose $q_{n+1}(\overline{y}, \overline{z})$ over the enlargement

     $$(C\overline{a}_{\leq n}\overline{e}^{\leq n}\overline{d}^{\leq n}\overline{a}_{n+1}; \mathcal{M}'_{n+1})$$

     extending q_n which does not þ-fork over $(C\overline{a}_{<n}\overline{e}^{<n}\overline{d}^{<n}; \mathcal{M}_n)$, and choose $\mathcal{M}_{n+1} \in K_{C\overline{a}_{\leq n}\overline{e}^{\leq n}\overline{d}^{\leq n}\overline{a}_{n+1}}$ and tuples $\overline{e}^{n+1}\overline{d}^{n+1}$ in $\mathcal{M}_{n+1}$ to realize q_{n+1}.

Finally, we may select a coherent sequence $(\mathcal{N}_n, (\overline{a}_n, \overline{e}^n, \overline{d}^n))_{n<\omega}$ patterned strictly on the sequence $(\mathcal{M}'_n, (\overline{a}_n, \overline{e}^n, \overline{d}^n))_{n<\omega}$.

Claim. *Let $m < n < \omega$. Then $\overline{a}_{m+1}...\overline{a}_n \mathop{\smile\hskip -0.9em ^| \ }^{þ}_{C\overline{d}^m} \overline{e}^m$.*

proof of claim. Given $m < \omega$, the claim is proved by induction on the difference $t = n - m$. The argument is identical to that in [22], so we omit it in the interest of brevity. $\square$

Now, up to some padding, we may assume that all of the $\overline{e}_i$'s have the same length, and all of the $\overline{d}_i$'s have the same length. Let $\overline{w}$ be a fresh tuple of variables, and define

$$\psi(\overline{y}, \overline{v}_1 \ldots \overline{v}_s \overline{w}) = \bigvee_{i=1}^{s} (\overline{w} = \overline{v}_i \wedge \varphi_i(\overline{y}, \overline{w}))$$

Furthermore, let

$$\delta_i(\overline{v}_i, \overline{z}, \overline{u}) = qtp(\overline{e}_i, \overline{d}_i, \overline{c}; \mathcal{M}_0)$$

where $\overline{c}$ is an enumeration of C, and let

$$\delta_i'(\overline{v}_1 \ldots \overline{v}_s \overline{w}, \overline{u}) = (\overline{w} = \overline{v}_i \wedge \delta_i(\overline{w}, \overline{z}, \overline{u}))$$

and $\Delta' = \{\delta_i' : i \in [s]\}$. Finally, let $r = \max_{i \in [s]} r_i$, where for each $i \in [s]$, the family $\left\{\varphi_i(\overline{y}, \overline{e}) : \overline{e} \vDash \delta_i(\overline{v}_i, \overline{d}_i)\right\}$ is setwise unbounded and r_i-inconsistent over $K_{C\overline{d}_i}$

Claim. *Let $m < n < \omega$. Then*

$$\text{Þ}(qtp(\overline{b}/C\overline{a}_m \ldots \overline{a}_n; \mathcal{N}_n), \psi, \Delta', r) < \text{Þ}(qtp(\overline{b}/C\overline{a}_{m+1} \ldots \overline{a}_n; \mathcal{N}_n), \psi, \Delta', r)$$

proof of claim. The argument is almost identical to the argument for the similar claim in the proof of proposition 5.6, so we refrain from duplicating it here. $\square$

Thus, by rank-monotonicity, $\text{þ}(qtp(\overline{b}/C; \mathcal{M}_0), \psi, \Delta', r) \geq \alpha + 1$ for all $\alpha < \omega$, which contradicts the assumption that K is rosy, and this completes the proof of the theorem. $\square$

Theorem 5.9 (Transitivity)**.** *Assume K is rosy. Let $M_0 \in K$, $B_1, B_2, C \subseteq M_0$ and $\overline{a}, \overline{b}_1, \overline{b}_2 \in M_0^{\leq k}$. Then, the following are equivalent:*

1. $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} \overline{b}_1 \overline{b}_2$
2. $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B_1$ *and* $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C\overline{b}_1} \overline{b}_2$.

Proof. We have already proved that 1⇒2 holds in all theories, so we only need to show here that 2⇒1. Let $M_0 \in K$, $B_1, B_2, C \subseteq M_0$ and $\overline{a} \in M_0^{\leq k}$, and suppose $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B_1$ and $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{B_1 C} B_2$. Then,

$$\begin{aligned}
\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} \overline{b}_1 \,\&\, \overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C\overline{b}_1} \overline{b}_2 &\implies \overline{b}_1 \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} \overline{a} \,\&\, \overline{b}_2 \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C\overline{b}_1} \overline{a} \\
&\implies \overline{b}_1 \overline{b}_2 \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} \overline{a} \\
&\implies \overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{C} B_1 B_2
\end{aligned}$$

by Symmetry;Partial left-transitivity;Symmetry, as required. $\square$

Finally, we can complete the characterization of þ-forking in terms of local þ-rank, theorem 5.6. Again, the proof is not significantly different from that in [22].

Proposition 5.10. *Assume K is rosy. Let $\mathcal{M} \in K$ and $B, C \subseteq M$, and let $p(\overline{x})$ be a complete quantifier-free type over $(BC; \mathcal{M})$. Suppose $p(\overline{x})$ does not þ-fork over C. Then $þ(p, \varphi, \Delta, r) = þ(p\restriction C, \varphi, \Delta, r)$ for all $1 < r < \omega$, $\varphi(\overline{x}, \overline{y})$ and Δ.*

Proof. We will prove, by induction on $\alpha < \omega$, that $þ(p\restriction C, \varphi, \Delta, r) \geq \alpha$ implies $þ(p, \varphi, \Delta, r) \geq \alpha$ for all p, $\varphi(\overline{x}, \overline{y})$, Δ and r as in the statement of the proposition (in particular, that p does not þ-fork over C). When $\alpha = 0$, the claim is obviously true. Assuming the claim for all $\alpha_0 \leq \alpha$, suppose that $þ(p\restriction C, \varphi, \Delta, r) \geq \alpha + 1$. In particular, let $\delta(\overline{y}, \overline{z}) \in \Delta$, $\mathcal{M}_1 \in K_{BC}$ and $\overline{d} \in M_1^{\overline{z}}$ be such that $\left\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash \delta(\overline{y}, \overline{d})\right\}$ is setwise unbounded and r-inconsistent over $K_{BC\overline{d}}$ and

$$\left\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash \delta(\overline{y}, \overline{d}),\ þ(p\restriction C \cup \{\varphi(\overline{x}, \overline{b})\}, \varphi, \Delta, r) \geq \alpha\right\}$$

is setwise unbounded over $K_{BC\overline{d}}$.

Claim. *Let $q(\overline{y})$ be a complete type over C which is K-consistent with $\delta(\overline{y}, \overline{d})$. Suppose $\mathcal{N} \in K_{BC\overline{d}}$ and $\overline{b} \in q(N^{\overline{y}})$ are such that $þ(p\restriction C \cup \{\varphi(\overline{x}, \overline{b})\}, \varphi, \Delta, r) \geq \alpha$. Then $þ(p \cup \{\varphi(\overline{x}, \overline{b})\}, \varphi, \Delta, r) \geq \alpha$.*

Proof. Fix $\overline{a} \vDash p$ – $\overline{a} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} B$ by hypothesis. By Extension, let $q'(\overline{y})$ be a complete extension of $q(\overline{y})$ to $BC\overline{a}$ which does not þ-fork over $C\overline{a}$, and without loss of generality, assume $\overline{b} \vDash q'$; thus, $\overline{b} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{a}} B$. Now, we derive:

$$\begin{aligned} \overline{a} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} B \,\&\, \overline{b} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{a}} B &\overset{\text{(Sym.)}}{\Longrightarrow} B \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} \overline{a} \,\&\, B \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{a}} \overline{b} \\ &\overset{\text{(Trans.)}}{\Longrightarrow} B \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} \overline{a}\overline{b} \\ &\overset{\text{(Trans.)}}{\Longrightarrow} B \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{b}} \overline{a} \\ &\overset{\text{(Sym.)}}{\Longrightarrow} \overline{a} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{b}} B \end{aligned}$$

Thus, $\overline{a} \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C\overline{b}} B$, and $þ(p\restriction C \cup \{\varphi(\overline{x}, \overline{b})\}, \varphi, \Delta, r) \geq \alpha$ by the assumptions of the claim. By the inductive hypothesis, $þ(p \cup \{\varphi(\overline{x}, \overline{b})\}, \varphi, \Delta, r) \geq \alpha$, which proves the claim. □

From the claim, we find that

$$\left\{\varphi(\overline{x}, \overline{b}) : \overline{b} \vDash \delta(\overline{y}, \overline{d}),\ þ(p \cup \{\varphi(\overline{x}, \overline{b})\}, \varphi, \Delta, r) \geq \alpha\right\}$$

is setwise unbounded over $K_{BC\overline{d}}$. Thus, $þ(p, \varphi, \Delta, r) \geq \alpha + 1$, as required. □

5.5 Sufficiency of symmetry and transitivity

In this section, we prove (theorem 5.13) that for K (with AP/sets) to be a rosy class, it is enough to verify that $\mathop{\smile\!\!\!\!\!\mid}^{\text{þ}}$ is symmetric of K. We present two technical lemmas leading up to the demonstration of theorem 5.13. The first of these is completely obvious:

Lemma 5.11. *There is a function $f_{ind} : \omega \times \omega \to \omega$ such that for any $\mathcal{M} \in K$, $A \subseteq M$, and $\overline{b}_1, ..., \overline{b}_m \in M^n$, if $m > f_{ind}(n, |A|)$ and $(\overline{b}_i)_{i=1}^m$ is A-indiscernible, then there is an infinite A-indiscernible coherent sequence $(\mathcal{M}_i, \overline{b}_i')_{i<\omega}$ in K_A such that*

$$qtp(\overline{b}_{i_1}'...\overline{b}_{i_m}'/A; \mathcal{M}_j) = qtp(\overline{b}_1...\overline{b}_m/A; \mathcal{M})$$

whenever $i_1 < \cdots < i_m \leq j < \omega$.

Lemma 5.12. *Let $\varphi(\overline{x}, \overline{y})$ and $\delta(\overline{y}, \overline{z})$ be quantifier-free formulas, and let $1 < r < \omega$. Suppose $þ((\overline{x} = \overline{x}), \varphi, \delta, r) = \infty$. Then there are $\mathcal{M} \in K$, $\overline{a} \in M^{\overline{x}}$, and an infinite $\overline{a}$-indiscernible coherent sequence $(\mathcal{M}_i, \overline{b}_i)_{i<\omega}$ such that for each $i < \omega$,*

1. *$\mathcal{M}_i \vDash \varphi(\overline{a}, \overline{b}_i)$*
2. *$\varphi(\overline{x}, \overline{b}_i)$ þ-divides over $\bigcup \{rng(\overline{b}_j) : j < i\}$*
3. *$qtp(\overline{b}_i/\overline{a}, \bigcup_{j<i} rng(\overline{b}_j); \mathcal{M}_i)$ does not þ-fork over $\bigcup_{j<i} rng(\overline{b}_j)$.*

Proof. For each $n < \omega$, we must implement the following inductive construction:

- Set $\pi_0(\overline{x}) = (\overline{x} = \overline{x}) = \bigwedge_i x_i = x_i$. Choose $\mathcal{M}_{n,0} \in K$, $\overline{a}_{n,0} \in M_{n,0}^{\overline{x}}$, $\overline{b}_{n,0} \in M_{n,0}^{\overline{y}}$ and $\overline{c}_{n,0} \in M_{n,0}^{\overline{z}}$ such that
  - $\mathcal{M}_{n,0} \vDash \varphi(\overline{a}_{n,0}, \overline{b}_{n,0}) \wedge \delta(\overline{b}_{n,0}, \overline{c}_{n,0})$;
  - $qtp(\overline{b}_{n,0}/\overline{c}_{n,0}; \mathcal{M}_{n,0})$ is non-algebraic
  - $þ(\pi_0(\overline{x}) \cup \{\varphi(\overline{x}, \overline{b}_{n,0})\}, \varphi, \delta, r) \geq n$
  - $\left\{\varphi(\overline{x}, \overline{b}') : \overline{b}' \vDash \delta(\overline{x}, \overline{c}_{n,0}),\ þ(\pi_0(\overline{x}) \cup \{\varphi(\overline{x}, \overline{b}')\}, \varphi, \delta, r) \geq n\right\}$ is setwise unbounded over $K_{\overline{c}_{n,0}}$
  - $\left\{\varphi(\overline{x}, \overline{b}') : \overline{b}' \vDash \delta(\overline{x}, \overline{c}_{n,0})\right\}$ is r-inconsistent over $K_{\overline{c}_{n,0}}$.

- Suppose we are given a finite coherent sequence $(\mathcal{M}_{n,j}, \overline{a}_{n,j}, \overline{b}_{n,j}, \overline{c}_{n,j})_{i \leq j}$ and types $\pi_0(\overline{x}) \subset \pi_1(\overline{x}) \subset \cdots \subset \pi_i(\overline{x})$, where π_t is over the set $D_{n,t} = \bigcup_{j<t} rng(\overline{b}_{n,j})$, and

  $$þ(\pi_t(\overline{x}), \varphi, \delta, r) \geq n - t + 1$$

  for each $t \leq i$.

  Then, we choose $\mathcal{M}_{n,i+1} \in K_{D_{n,i}}$, $\overline{a}_{n,i+1} \in M_{n,i+1}^{\overline{x}}$, $\overline{b}_{n,i+1} \in M_{n,i+1}^{\overline{y}}$ and $\overline{c}_{n,i+1} \in M_{n,i+1}^{\overline{z}}$ such that

– $\mathcal{M}_{n,i+1} \models \varphi(\overline{a}_{n,i+1}, \overline{b}_{n,i+1}) \wedge \delta(\overline{b}_{n,i+1}, \overline{c}_{n,i+1})$;

– $qtp(\overline{b}_{n,i+1}/D_{n,i}\overline{c}_{n,i+1}; \mathcal{M}_{n,i+1})$ is non-algebraic

– $qtp(\overline{b}_{n,i+1}/D_{n,i} \cup \bigcup_{j\leq i} rng(\overline{a}_{n,j}); \mathcal{M}_{n,i+1})$ does not þ-fork over $D_{n,i}$

– $þ(\pi_i(\overline{x}) \cup \{\varphi(\overline{x}, \overline{b}_{n,i+1})\}, \varphi, \delta, r) \geq n$

– $\left\{\varphi(\overline{x}, \overline{b}') : \overline{b}' \models \delta(\overline{x}, \overline{c}_{n,i+1}),\ þ(\pi_i(\overline{x}) \cup \{\varphi(\overline{x}, \overline{b}')\}, \varphi, \delta, r) \geq n\right\}$ is setwise unbounded over $K_{\overline{c}_{n,0}}$

– $\left\{\varphi(\overline{x}, \overline{b}') : \overline{b}' \models \delta(\overline{x}, \overline{c}_{n,0})\right\}$ is r-inconsistent over $K_{\overline{c}_{n,0}}$.

In this manner, we obtain a sequence of finite coherent sequences

$$\Gamma = \Big(\left(\mathcal{M}_{n,i}, \overline{a}_{n,i}, \overline{b}_{n,i}, \overline{c}_{n,i}\right)_{i\leq n} \Big)_{n<\omega}$$

Applying König's lemma, we extract from Γ a single infinite coherent sequence $(\mathcal{M}_n, \overline{a}_n\overline{b}_n)_{n<\omega}$ such that for all $m \leq n < \omega$, $\mathcal{M}_n \models \varphi(\overline{a}_n, \overline{b}_m)$ and $\varphi(\overline{x}, \overline{b}_n)$ þ-divides over $C_{<n} = \bigcup_{i<n} rng(\overline{b}_i)$. Applying Ramsey's theorem, we may assume that $(\mathcal{M}_n, \overline{a}_n, \overline{b}_n)_{n<\omega}$ is indiscernible.

Now, suppose $f_{\text{ind}}(|\overline{x}|, |\overline{y}|) < m < \omega$, and again applying Ramsey's theorem, choose $s(m) < \omega$ large enough to ensure that if $n \geq s(m)$, then $(\overline{b}_i)_{i<n}$ contains an $\overline{a}_n$-indiscernible subsequence of length at least m. In $\mathcal{M}_n$, let $i_1 < \cdots i_m < n$ be such that $(\overline{b}_{i_j})_{j=1}^m$ is $(rng(\overline{a}_n); \mathcal{M}_n)$-indiscernible. Note that by definition of þ-dividing, $\varphi(\overline{x}, \overline{b}_{i_t})$ þ-divides over $\bigcup_{u<t} rng(\overline{b}_{i_u})$ for each $t \leq m$.

Up to renaming elements, we fix a tuple $\overline{a}$, and using the argument of the previous paragraph, we obtain a sequence $\Gamma^* = (\mathcal{M}_n, (\overline{b}_{n,i})_{i\leq n})_{n<\omega}$ over $K_{\overline{a}}$ (not necessarily coherent) such that $\mathcal{M}_n \models \varphi(\overline{a}, \overline{b}_{n,i})$ for all $i \leq n < \omega$ and $\varphi(\overline{x}, \overline{b}_{n,i})$ þ-divides over $\bigcup_{j<i} rng(\overline{b}_{n,j})$ for each $i \leq n$. By applying techniques we've deployed several times previously, we then extract an infinite $\overline{a}$-indiscernible coherent sequence strictly patterned on Γ^* – and by definition of Γ^*, it has the desired properties. □

Hopefully, the proof of preceding lemma, 5.12, adequately demonstrates how gnarly arguments that completely avoid the Compactness theorem can be. The next theorem can also be proved without appeal to the Compactness theorem, although this proof involves, essentially, duplicating the preceding argument while satisfying some additional constraints – hence, demanding many more applications of König's lemma and Ramsey's theorem. To save ourselves and the reader this horror, we give, together with the bald assertion that it is not strictly necessary, an argument that *does* involve the Compactness theorem.

Theorem 5.13. *Assume K has an $\aleph_0$-categorical K-universal limit model, $\mathfrak{M}$ (a Fraŝse limit). If $\mathop{\smile\!\!\!\!\!\!|}^{þ}$ is symmetric in K, then K is a rosy class.*

Proof. We prove the contrapositive assertion. Suppose K is not rosy – in particular, we may suppose that there are $\varphi(\overline{x}, \overline{y})$ and $\delta(\overline{y}, \overline{z})$, quantifier-free formulas without parameters, and some $1 < r < \omega$ such that $þ(\overline{x} = \overline{x}, \varphi, \delta, r) = \infty$. By lemma 5.12, there are $\mathcal{M} \in K$, $\overline{a} \in M^{\overline{x}}$, and an infinite $\overline{a}$-indiscernible coherent sequence $(\mathcal{M}_n, \overline{b}_n)_{n<\omega}$ such that for each $n < \omega$, $\mathcal{M}_n \models \varphi(\overline{a}, \overline{b}_n)$, $\varphi(\overline{x}, \overline{b}_n)$ þ-divides over $\bigcup_{i<n} rng(\overline{b}_i)$ and $qtp(\overline{b}_n/\overline{a}, \bigcup_{i<n} rng(\overline{b}_i); \mathcal{M}_n)$ does not þ-fork over $\bigcup_{i<n} rng(\overline{b}_i)$. Now, the next claim follows directly from the Compactness theorem and the fact that $\mathfrak{M}$ is $\aleph_0$-categorical.

Claim. *There are $\overline{a} \in ||\mathfrak{M}||^{\overline{x}}$ and an $\overline{a}$-indiscernible sequence $(\overline{b}_n)_{n\leq\omega}$ in $||\mathfrak{M}||^{\overline{y}}$ such that*

- $\mathfrak{M} \models \varphi(\overline{a}, \overline{b}_n)$ *for all $n \leq \omega$*
- $\varphi(\overline{x}, \overline{b}_n)$ *þ-divides over $\bigcup_{i<n} rng(\overline{b}_i)$ for all $n \leq \omega$*
- $qtp(\overline{b}_n/\overline{a}, \bigcup_{i<n} rng(\overline{b}_i); \mathfrak{M})$ *does not þ-fork over $\bigcup_{i<n} rng(\overline{b}_i)$.*

From the claim, it then follows immediatly that $\overline{b}_\omega \mathop{\smile\!\!\!\!\!\!|}^{þ}_{\overline{b}_0} \overline{a}$. Furthermore,

$$þ(qtp(\overline{a}/\overline{b}_0; \mathfrak{M}), \varphi, \delta, r) \geq þ(qtp(\overline{a}/\overline{b}_0\overline{b}_\omega; \mathfrak{M}), \varphi, \delta, r) + 1$$

by the construction, so $\overline{a} \not\mathop{\smile\!\!\!\!\!\!|}^{þ}_{\overline{b}_0} \overline{b}_\omega$. Thus, the failure of rosiness implies the failure of the symmetry of $\mathop{\smile\!\!\!\!\!\!|}^{þ}$, as desired. $\square$

Finally, we extend the relation $\mathop{\smile\!\!\!\!\!\!|}^{þ}$ to triples of (finite) sets as follows: If $\mathcal{M} \in K$, $A, B, C \subseteq M$, then

$$A \mathop{\smile\!\!\!\!\!\!|}^{þ}_C B \;\Leftrightarrow\; \overline{a} \mathop{\smile\!\!\!\!\!\!|}^{þ}_C B \text{ whenever } \overline{a} \in A^k$$

It is, then, not difficult to show (from Symmetry) that, provided K is rosy,

$$A \mathop{\smile\!\!\!\!\!\!|}^{þ}_C B \;\Leftrightarrow\; \overline{a} \mathop{\smile\!\!\!\!\!\!|}^{þ}_C \overline{b} \text{ whenever } \overline{a} \in A^k \text{ and } \overline{b} \in B^k$$

Thus, we have the principal theorem of this chapter:

Theorem 5.14. *Let T be the complete k-variable theory of a finite ρ-structure which is capped and has infinitely many finite models up to isomorphism, and let $K =$ **fin**$[T^G]$. Then, K is rosy if and only if $\mathop{\smile\!\!\!\!\!\!|}^{þ}$ is a true independence relation.*

5.6 Imaginaries, and why we haven't talked about them until now

In this section, we will introduce and discuss the notion of "imaginary elements" that seems to be most natural (in our opinion) for a class $K = K^G$ of finite structures with amalgamation over sets. To remove any suspense for the treatment, the key observation –

towards retaining our analysis of þ-independence *mutatis mutandis* – is simply that only finitely many imaginary sorts carry any "new" information with respect to a finite base set.

Obviously, we proceed with the definitions first. Let K be a class of finite ρ-structures with amalgamation over sets. Let $\theta = \theta(x_1, ..., x_m, y_1, ..., y_n)$ be quantifier-free formula in the language of ρ (with an explicit separation of variables), and let $\mathcal{M} \in K$. For convenience in the sequel, we assume that ρ contains a unary predicate symbol S_0 such that $S_0^{\mathcal{N}} = N$ for all $\mathcal{N} \in K$ We expand the signature ρ to t $\rho_\theta = \rho \cup \{S_\theta^{(1)}, P_\theta^{(n+1)}\}$, and for $\mathcal{M} \in K$, we define a ρ_θ-expansion $\mathcal{M}^\theta$ with new interpretations as follows:

- $S_\theta^{\mathcal{M}^\theta} = \left\{\theta(M^m, \bar{b}) : \bar{b} \in M^n\right\} \setminus \{\emptyset\}$
- $||\mathcal{M}^\theta|| = M \,\dot{\cup}\, S_\theta^{\mathcal{M}^\theta}$
- $P_\theta^{\mathcal{M}^\theta} = \left\{(\bar{b}, \theta(M^m, \bar{b})) : \bar{b} \in M^n\right\}$

Further, if $A \subseteq M$, we also define a structure $(A; \mathcal{M})^\theta$ as follows:

- $D_\theta(A; \mathcal{M}) = \left\{\bar{b} \in M^n : A^m \cap \theta(M^m, \bar{b}) \neq \emptyset\right\}$
- $S_0^{(A;\mathcal{M})^\theta} = A$
- $S_\theta^{(A;\mathcal{M})^\theta} = \left\{A^m \cap \theta(N^m, \bar{b}) : \mathcal{N} \in K_A,\, \bar{b} \in D_\theta(A; \mathcal{N})\right\} \setminus \{\emptyset\}$
- $||(A; \mathcal{M})^\theta|| = S_0^{(A;\mathcal{M})^\theta} \,\dot{\cup}\, S_\theta^{(A;\mathcal{M})^\theta}$
- $P_\theta^{(A;\mathcal{M})^\theta} = \left\{(\bar{b}, A^m \cap \theta(M^m, \bar{b})) : \bar{b} \in D_\theta(A; \mathcal{M}) \cap A^n\right\}$
- $R^{(A;\mathcal{M})^\theta} = A^r \cap R^{\mathcal{M}}$ whenever $R \in \rho$ and $r = ari(R)$.

Observation. *Let θ_1 and θ_2 be quantifier-free formulas a in the language of ρ (with an explicit separation of variables). Suppose $\mathcal{M} \in K$ and $A \subseteq M$. Then $(\mathcal{M}^{\theta_1})^{\theta_2} = (\mathcal{M}^{\theta_2})^{\theta_1}$ and $(\mathcal{M}^{\theta_1})^{\theta_2} = (\mathcal{M}^{\theta_2})^{\theta_1}$.*

It can also be shown that if θ_1 and θ_2 are both K-adherent (defined below), then θ_2 is K^{θ_2}-adherent.

Thus, if Θ is a set of quantifier-free formulas in the language of ρ (with explicit separations of variables), then the operators $\mathcal{M} \mapsto \mathcal{M}^\Theta$ and $(A; \mathcal{M}) \mapsto (A; \mathcal{M})^\Theta$, specified via any linear order of Θ, are well-defined.

Our first two lemmas of this section (whose fairly routine proofs we omit) demonstrate that there is – as in the full first-order scenario –essentially only one "level" in construction of imaginaries; more precisely, every imaginary element is determined by its trace on the real universe.

Lemma 5.15. *Let $\psi(\overline{x}, \overline{x}'; \overline{y})$ be a quantifier-free formula of ρ_Θ, for some set Θ of quantifier-free formulas of the language of ρ, and assume that*

$$\psi \models \bigwedge_i S_0(x_i) \wedge \bigwedge_j \neg S_0(x'_j)$$

Then there is a quantifier-free formula $\hat{\psi}(\overline{x}''; \overline{y})$ such that $\hat{\psi} \models \bigwedge_l S_0(x''_l)$ and for every $\mathcal{M} \in K$ and all $\overline{b}_1, \overline{b}_2 \in ||\mathcal{M}^\Theta||^{\overline{y}}$,

$$\hat{\psi}(M^{\overline{x}''}, \overline{b}_1) = \hat{\psi}(M^{\overline{x}''}, \overline{b}_2) \;\Leftrightarrow\; \psi(M^{\overline{x}}, ||\mathcal{M}^\Theta||^{\overline{x}'}, \overline{b}_1) = \psi(M^{\overline{x}}, ||\mathcal{M}^\Theta||^{\overline{x}'}, \overline{b}_2)$$

We will now specify two criteria which we believe are both natural and intrinsic to the notion of imaginary elements (in the full first-order scenario) as the the notion is "pushed down" onto our scenario in this dissertation. The first criterion is the analog of what happens in the full first-order case when considering a common diagram of structures $\mathcal{M}_1^{\mathrm{eq}}$ and $\mathcal{M}_2^{\mathrm{eq}}$. The second criterion is the analog of the observation that (up to transformation by a canonical functor) adding imaginaries does not affect the theory in question.

K**-adherence**. Let $\theta(\overline{x}, \overline{y})$ be a quantifier-free formula

1. Let $\mathcal{M}, \mathcal{N} \in K$, let $f : \mathcal{M} \to \mathcal{N}$ be an embedding. Then there is a unique *good expansion* $f^\theta : \mathcal{M}^\theta \to \mathcal{N}^\theta$ of f determined by the rule

   $$f^\theta(\,\theta(M^{\overline{x}}, \overline{b})\,) = \theta(N^{\overline{x}}, f\overline{b})$$

   If $A \subseteq M$, then there are such good expansions $f'(A; \mathcal{M})^\theta \to \mathcal{N}^\theta$, but they need no be unique. Now, consider the situation of a sequence of embeddings:

   $$(A; \mathcal{M}) \xrightarrow{f} \mathcal{N} \xrightarrow{g} \mathcal{N}_1$$

   We then say that the expansion f' is g-compatible just in case $g^\theta \circ f'$ is a good expansion of $g \circ f$.

2. We say that $\theta(\overline{x}, \overline{y})$ is K-*adherent* if for any amalgamation quartet of K

   $$\begin{tikzcd} (A; \mathcal{M}_0) \arrow[r, "f_1"] \arrow[d, "f_2"'] & \mathcal{M}_1 \arrow[d, "g_1"] \\ \mathcal{M}_2 \arrow[r, "g_2"'] & \mathcal{N} \end{tikzcd}$$

   for any good expansions f'_1, f'_2 of f_1, f_2, respectively, f'_i is g_i-compatible, $i = 1, 2$, and $g_1^\theta \circ f'_1 = g_2^\theta \circ f'_2$.

T-admissibility. Again, let $\theta(\overline{x},\overline{y})$ be a quantifier-free formula of ρ. We say that $\theta(\overline{x},\overline{y})$ is T-*admissible* if $|\overline{y}| \leq k$ and $\mathcal{M}^\theta \equiv^k \mathcal{N}^\theta$ for all $\mathcal{M},\mathcal{N} \in K = \mathbf{fin}[T^G]$

The following proposition is an easy consequence of the fact that everything in sight is finite.

Proposition 5.16. *Let $\theta_1(\overline{x}_1,\overline{y})$ and $\theta_2(\overline{x}_2,\overline{y})$ be K^G-adherent quantifier-free formulas, and assume that $\theta_1(\overline{x}_1,\overline{y})$ is T-admissible. Suppose that*

($\star$) for some $\mathcal{M} \in K^G$, for all $\overline{b},\overline{b} \in M^{\overline{y}}$, if $qtp(\overline{b};\mathcal{M}) = qtp(\overline{c};\mathcal{M})$, then

$$\theta_1(M^{\overline{x}_1},\overline{b}) = \theta_1(M^{\overline{x}_1},\overline{c}) \;\Leftrightarrow\; \theta_2(M^{\overline{x}_2},\overline{b}) = \theta_2(M^{\overline{x}_2},\overline{c})$$

Then, the following are equivalent:

1. *$\theta_2(\overline{x}_2,\overline{y})$ is T-admissible;*
2. *$T \vDash \forall\overline{y},\overline{z}\,[\,\forall\overline{x}_1(\theta(\overline{x}_1,\overline{y}) \leftrightarrow \theta_1(\overline{x}_1,\overline{z})) \leftrightarrow \forall\overline{x}_2(\theta(\overline{x}_2,\overline{y}) \leftrightarrow \theta_2(\overline{x}_2,\overline{z}))\,]$*

If θ_1 and θ_2 satisfy the condition $\star$ in proposition 5.16, then we say that θ_1 *subsumes* θ_2 *over* K. Let $Q(T)$ denote the set of all quantifier-free formulas $\theta(\overline{x},\overline{y})$ which are both K-adherent and T-admissible, and let $\prec$ be any linear order of $Q(T)$. Then, define $\pi_\prec : Q(T) \to Q(T)$ so that for any $\theta(\overline{x},\overline{y})$, up to substitutions of variables, $\theta'(\overline{x}',\overline{y}) = \pi_\prec(\theta(\overline{x},\overline{y}))$ is the $\prec$-least formula that subsumes θ over K.

Observation. *$\Phi_K(\prec) = img(\pi_\prec)$ is finite.*

Now, consider $\mathcal{M} \in K$ and its expansion $\mathcal{M}^{\Phi_K(\prec)}$, which necessarily has either infinitely many or no sorts beyond ρ^G. Let

$$\tau^{(\prec)} = \rho^G \cup \{S_\theta : \theta \in \Phi_K(\prec)\} \cup \{P_\theta : \theta \in Q(T)\}$$

be a new signature. We obtain a $\tau^{(\prec)}$-expansion $\mathcal{M}^{(\prec)}$ (of $\mathcal{M}$) by interpreting the new relation symbols as follows:

1. $S_\theta^{\mathcal{M}^{(\prec)}} = S_\theta^{\mathcal{M}^\theta}$ whenever $\theta \in \Phi_K(\prec)$
2. Let $\theta(\overline{x},\overline{y}) \in Q(T)$ and $\theta_0(\overline{x}_0,\overline{y}) = \pi_\prec(\theta)$; then set

$$P_\theta^{\mathcal{M}^{(\prec)}} = \left\{(\overline{b},\theta_0(M^{\overline{x}_0},\overline{b})) : \overline{b} \in M^{\overline{y}}, \theta(M^{\overline{x}},\overline{b}) \neq \emptyset\right\}$$

Let $\prec_1$ and $\prec_2$ be linear orders of $Q(T)$. It is quite easy to see that, up to renaming the relation symbols naming the imaginary sorts, $\mathcal{M}^{(\prec_1)}$ and $\mathcal{M}^{(\prec_2)}$ are identical. Thus, specifying a linear order $\prec$ of $Q(T)$ arbitrarily, we define $\rho^{\mathrm{eq}} = \tau^{(\prec)}$, $\mathcal{M}^{\mathrm{eq}} = \mathcal{M}^{(\prec)}$ for $\mathcal{M} \in K$, and $K^{\mathrm{eq}} = \{\mathcal{M}^{\mathrm{eq}} : \mathcal{M} \in K\}$ without any significant ambiguity.

For brevity, we write M^{eq} as a shorthand for $||\mathcal{M}^{\mathrm{eq}}||$. Furthermore, if $\mathcal{M} \in K$, $A \subseteq M^{\mathrm{eq}}$ and $0 < n < \omega$, then we let $S_n(A;\mathcal{M}^{\mathrm{eq}})$ to denote the set of complete quantifier-free n-types over A in the language of ρ^{eq}

Observation. *The following are all easily verified.*

1. *If K admits amalgamation over sets, then K^{eq} does, too.*
2. *There is a function $h : \omega \times \omega \to \omega$ such that for any $\mathcal{M} \in K$, $A \subseteq M^{eq}$ and $0 < n < \omega$, $|S_n(A; \mathcal{M}^{eq})| \leq h(|A|, n)$.*
3. *There is a function $g_{rk} : \omega \times \omega \to \omega$ such that for any $n < \omega$, $\mathcal{N} \in K$, $B_0 \subseteq B \subseteq N$, and any complete type $p(\overline{x})$ over $(B; \mathcal{N})$, the following are equivalent:*
   (a) *$\text{þ}(p, \Phi, \Delta, r) < \text{þ}(p{\restriction}B_0, \Phi, \Delta, r)$ for some Φ, Δ and r.*
   (b) *With $m = g_{rk}(n, |B_0|)$,*

$$\text{þ}(p, G_m, G_m, m) < \text{þ}(p{\restriction}B_0, G_m, G_m, m)$$

where G_m is the set of $\leq$ m-ary quantifier-free ρ^{eq}-formulas with non-equality atoms in

$$\rho^G \cup \{S_\theta : \theta \in \Phi_K(\prec)\} \cup \{P_\theta : \theta = \theta(x_1, ..., x_l, \overline{y}),\ l \leq m\}$$

counted up to logical equivalence.

Our development of þ-independence and rosiness, can then be recovered *mutatis mutandis.*

Chapter 6

Characterizing rosy classes

This chapter contains a number of results that are essential to our analysis of efficient computability of the model-building problem but which don't fit very nicely into any of the other chapters. The first of these is a characterization of þ-independence and rosiness analogous to the Independence theorem for simple theories. More precisely, we show (following [10]) that $\mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}$ is the coarsest possible notion of independence on any class K, and if K admits any true independence relation at all, then K is rosy.

A reader with a background in model theory will probably have noticed that we have not yet made any mention of Local character thus far. As it turns out, Local character is inessential to the formulation of rosiness for classes of finite structures. Indeed the characterization mentioned in the previous paragraph does not require it, in contrast to the situation in the first-order model theory of infinite structures. In our situation, we find that the presence of Local character is equivalent to a stronger condition which might be called "super-rosiness."

6.1 Local character, $U^{\text{þ}}$-rank and small algebraicity

6.1.1 Local character

Up to this point, we have had no need of the property known as Local character in the model theory literature, and we have held off from this part of the development because establishing the definition of Local character is, to some degree, a philosophical matter. The definition we present certainly "works," but in contrast to the full first-order scenario, we don't have the cardinality of the theory to use as point of comparison. Further, the definition we give already enforces a condition that would be identified as "super-rosiness" in the literature. In any case, we say that $\mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}$ *has Local character over* K if there is a function $f_{\mathrm{loc}} : \omega \to \omega$ such that for all $\mathcal{M} \in K$ and $A, B \subseteq M$, $A \mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}_{B_0} B$ for some $B_0 \in \binom{B}{f_{\mathrm{loc}}(|A|)}$.

Proposition 6.1. *If* $\mathop{\smile\!\!\!\!\!\!|}^{\text{þ}}$ *has Local character over* K*, then* K *is rosy.*

Proof. Assume $\mathop{\smile\!\!\!\!\!\!\downarrow}^{\text{þ}}$ has Local character over K via $f_{\text{loc}} : \omega \to \omega$. For a contradiction, suppose $\varphi(\overline{x}, \overline{y})$, $\delta(\overline{y}, \overline{z})$ are quantifier-free formulas and $1 < r < \omega$ are such that $\text{þ}(\overline{x} = \overline{x}, \varphi, \delta, r) = \infty$. By lemma 5.12, there are $\mathcal{M}_0 \in K$, $\overline{a} \in M_0^{\overline{x}}$, and an $\overline{a}$-indiscernible coherent sequence $(\mathcal{M}_i, \overline{b}_i)_{i<\omega}$ such that

1. $\mathcal{M}_i \vDash \varphi(\overline{a}, \overline{b}_i)$
2. $\varphi(\overline{x}, \overline{b}_i)$ þ-divides over $B_i = \bigcup \{rng(\overline{b}_j) : j < i\}$

whenever $0 < i < \omega$. By Local character, there is a subset $C \subset \bigcup_{i<\omega} B_i$ such that $|C| \leq f_{\text{loc}}(|\overline{x}|)$ and $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{\text{þ}}_C B_i$ for all large enough $i < \omega$. (More precisely, for each $i < \omega$, there is a subset $C_i \subseteq B_i$ such that $|C_i| \leq f_{\text{loc}}(|\overline{x}|)$ and $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{\text{þ}}_{C_i} B_i$. By monotonicity, $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{\text{þ}}_{C_i} B_j$ whenever $j \leq i$, so we may assume that $C_j \subseteq C_i$ whenever $j \leq i < \omega$.) This contradicts the fact that $\varphi(\overline{x}, \overline{b}_i) \in qtp(\overline{a}/B_{i+1}; \mathcal{M}_{i+1})$ þ-divides over B_i. Thus, K is rosy. $\square$

6.1.2 $U^{\text{þ}}$-rank and small algebraicity

We use a standard definition of the $U^{\text{þ}}$-rank, adapted slightly for our purposes: Let $\mathcal{M} \in K$, $C \subseteq M_0$ and $p(\overline{x}) \in S_k^{\text{qf}}(C; \mathcal{M}_0)$

- $U^{\text{þ}}(p) \geq 0$
- $U^{\text{þ}}(p) \geq \alpha + 1$ if there are $\mathcal{M} \in K_C$, $C \subseteq D \subseteq M$ and $p'(\overline{x}) \in S_k^{\text{qf}}(D; \mathcal{M})$ extending $p(\overline{x})$ such that (i) p' þ-forks over $(C; \mathcal{M})$ and (ii) $U^{\text{þ}}(p') \geq \alpha$.
- $U^{\text{þ}}(p) = \infty$ just in case $U^{\text{þ}}(p) \geq \alpha$ for all $\alpha < \omega$.

If $p(\overline{x}) = qtp(\overline{a}/C; \mathcal{M})$, then obviously we may write $U^{\text{þ}}(\overline{a}/C; \mathcal{M})$ in place of $U^{\text{þ}}(p(\overline{x}))$ without ambiguity.

The first goal of this section is to find a natural condition (with respect to efficient computability) which is enough to ensure that in a rosy class K satisfying that condition, $U^{\text{þ}}$ is always defined and takes only finite values. The condition we isolate – *small algebraicity* – was shown in chapter 4 to obtain whenever the model-building problem for K admits a (relational) polynomial-time solution.

Thus, we assume that K is an $\forall\exists$-axiomatized class of finite ρ-structures with amalgamation over sets. Suppose $\mathfrak{M}$ is a countable Fraïssé limit of K; in particular, $\mathfrak{M}$ is ultrahomogeneous and $Th(\mathfrak{M})$ is $\aleph_0$-categorical and eliminates quantifiers. Recall, then, that K has small algebraicity if there is a number $d_K < \omega$ such that for all $C \subset_{\mathbf{fin}} ||\mathfrak{M}||$, if $p(x) \in S_1^{\text{qf}}(C; \mathfrak{M})$ is algebraic, then $p(x){\restriction}C_0$ is already algebraic for some $C_0 \in \binom{C}{d_K}$.

The first fact we prove seems to be incidental, but we include it anyway, for some reason. Subsequently, we derive characterizations of strong dividing and þ-dividing that will allow us to prove that Local character holds under the assumption of small algebraicity.

Lemma 6.2. *Let K and $\mathfrak{M}$ be as above, and assume that K is rosy and has small algebraicity. Let $0 < t < \omega$. If $\overline{a} \in ||\mathfrak{M}||^t$ and*

$$C_0 \subset C_1 \subset \cdots \subset C_n \subset \cdots \subset_{\boldsymbol{fin}} M$$

is a strictly increasing chain of algebraically closed finite sets, then there is a number n_0

$$acl^{\mathfrak{M}}(\overline{a}C_n) = acl^{\mathfrak{M}}(\overline{a}C_{n_0}) \cup C_n$$

whenever $n_0 \leq n < \omega$.

Proof. Towards a contradiction, suppose there are $\overline{a} \in ||\mathfrak{M}||^t$, a chain of finite algebraically closed sets

$$C_0 \subset C_1 \subset \cdots \subset C_n \subset \cdots \subset_{\mathbf{fin}} M$$

and for each $n < \omega$

$$b_{n+1} \in acl^{\mathfrak{M}}(\overline{a}C_{n+1}) \setminus (acl^{\mathfrak{M}}(\overline{a}C_n) \cup C_{n+1})$$

By the assumption of small algebraicity, for each $n < \omega$, there is a tuple $\overline{c}_{n+1} \in C_{n+1}^{d_K}$ such that $qtp(b_{n+1}/\overline{a}\overline{c}_{n+1})$ is algebraic. Of course, for each $n < \omega$, $qtp(b_n/\overline{a})$ is non-algebraic.

By the pigeonhole principle, we may assume that $\varphi(\overline{x}, y) = qtp(\overline{a}, b_n)$ and $\delta(y, \overline{z}) = qtp(b_n/\overline{c}_n)$ are both a constant over $n < \omega$. With $r = 1 + f_{\text{alg}}(t + d_K)$, it is then routine to verify that $\text{þ}(qtp(\overline{a}), \varphi, \delta, r) = \infty$, contradicting the hypothesis that K is rosy. □

Lemma 6.3 (Triviality of strong dividing)**.** *Assume that K has small algebraicity. Let $\mathcal{M} \in K$, $C \subseteq M$ and $\overline{a}, \overline{b}, \overline{e} \in M^{<\omega}$, and let $p(\overline{x}, \overline{y}) = qtp(\overline{a}, \overline{b}/C; \mathcal{M})$. If $p(\overline{x}, \overline{b})$ divides strongly over $C\overline{e}$, then there are $b^* \in rng(\overline{b})$ and $D \in \binom{C\overline{e}}{d_K}$ such that $qtp(\overline{a}, b^*/D; \mathcal{M})$ divides strongly in $qtp(b^*/D)$.*

Proof. By definition of strong dividing, $qtp(\overline{b}/C\overline{e}\overline{a})$ is algebraic, while $qtp(\overline{b}/C\overline{a})$ is not. Since K is a Fraïssé class, we know that each $b \in rng(\overline{b})$ is also algebraic over $(C\overline{e}\overline{a}; \mathcal{M})$, so that $qtp(\overline{a}, b/C)$ divides strongly over $C\overline{e}$ in $qtp(b/C\overline{e})$ provided $qtp(b/C\overline{e})$ is not already algebraic. Thus, we choose any $b^* \in rng(\overline{b})$ such that $qtp(b^*/C\overline{e}\overline{a})$ is algebraic and $qtp(b^*/C\overline{e})$ is not. By the small algebraicity assumption, there is a subset $D \in \binom{C\overline{e}}{d_K}$ such that $qtp(b^*/D\overline{a})$ is algebraic. Of course, $qtp(b^*/D)$ is not algebraic, so $qtp(\overline{a}, b^*/D)$ divides strongly in $qtp(b^*/D)$. □

Proposition 6.4 (þ-Dividing configurations)**.** *Assume that K has small algebraicity. Let $\mathcal{M}_0 \in K$, $\overline{a} \in M_0^{<\omega}$ and $b \in M_0$, and let $p_0(\overline{x}, y) = qtp(\overline{a}, b; \mathcal{M})$. There is a (finite) set $\mathcal{D}(p_0)$ of types of the form $q(\overline{x}, y, z_1, ..., z_l)$ laterally extending p_0, $l \leq d_K$, such that if $\mathcal{M} \in K$, $C \subseteq M$, and $p(\overline{x}, y) \in S^{qf}(C; \mathcal{M})$ is an extension of p_0, then the following are equivalent:*

1. *$p(\overline{x}, b)$ þ-divides over $(C; \mathcal{M})$.*
2. *$qtp(b/C; \mathcal{M})$ is non-algebraic and for some type $q(\overline{x}, y, \overline{z}) \in \mathcal{D}(p_0)$ and $\overline{c} \in C^{\overline{z}}$, $q(\overline{x}, b, \overline{c}) \subseteq p(\overline{x}, b)$.*

Proof. This is an easy-enough consequence of the triviality of strong dividing. $\square$

Now, we define $\mathcal{A}(K)$ to be the set of types $q(y, z_1, ..., z_{d_K})$ such that if $\mathcal{M} \in K$, $b \in M$, $\bar{c} \in M^{\bar{z}}$, and $\mathcal{M} \models q(b, \bar{c})$, then $b \in acl(\bar{c}; \mathcal{M})$. Using the families $\mathcal{D}(-)$ and $\mathcal{A}(K)$, the next proposition follows almost immediately:

Proposition 6.5 (Local character)**.** *Assume that K is rosy with small algebraicity. Let $\mathcal{M}_0 \in K$, $\bar{a}_0, \bar{b}_0 \in M_0^k$, and let $p_0(\bar{x}, y) = qtp(\bar{a}_0, \bar{b}_0; \mathcal{M})$. Then there is a finite set $\mathbf{F}(p_0)$ of lateral extensions $q(\bar{x}, \bar{y}, \bar{z})$ of p_0 such that for any $\mathcal{M} \in K$, $C_0 \subseteq C \subseteq M$, $\bar{a}, \bar{b} \in M^k$ such that for $p_0 \subseteq p(\bar{x}, \bar{y}) = qtp(\bar{a}, \bar{b}/C; \mathcal{M})$, the following are equivalent:*

1. *$p(\bar{x}, \bar{b})$ does not þ-fork over C_0;*
2. *For each $q(\bar{x}, \bar{y}, \bar{z}) \in \mathbf{F}(p_0)$, if there is a $\bar{c} \in C^{\bar{z}}$ such that $\mathcal{M} \models q(\bar{a}, \bar{b}, \bar{c})$, then there is a $\bar{c}_0 \in C_0^{\bar{z}}$ such that $\mathcal{M} \models q(\bar{a}, \bar{b}, \bar{c}_0)$*

As a corollary (using symmetry of $\mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\vert$}}\nolimits^{þ}$), we determine that $\mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\vert$}}\nolimits^{þ}$ has Local character over K.

Theorem 6.6. *Assume that K is rosy with small algebraicity. Then $U^{þ}$-rank is defined and finite-valued for every type in K.*

Proof. Naturally, we suppose

$$C_0 \subset C_1 \subset \cdots \subset C_n \subset \cdots \subset_{\mathbf{fin}} ||\mathfrak{M}||$$

is an infinite ascending chain of finite subsets of $\mathfrak{M}$, where $\mathfrak{M}$ is the countable Fraïssé limit of K, and for each $n < \omega$, let $p_n(\bar{x})$ be a complete quantifier-free type over $(C_n; \mathfrak{M})$ so that $p_m(\bar{x}) \subset p_n(\bar{x})$ whenever $m < n < \omega$. By Local character, for each $n < \omega$, there is a subset $C_{n,0} \in \binom{C_n}{f_{\text{loc}}(|\bar{x}|)}$ such that p_n does not þ-fork over $C_{n,0}$. By monotonicity, there is a number $t < \omega$ such that if $n \geq t$, then p_n does not þ-fork over $D = \bigcup_{m \leq t} C_{m,0} \subseteq C_n$. In particular, if $n > t$ and $\bar{a} \models p_n$, then $\bar{a} \mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\vert$}}\nolimits^{þ}_D C_n$, which implies that $\bar{a} \mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\vert$}}\nolimits^{þ}_{C_{n-1}} C_n$ by the transitivity property of $\mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\vert$}}\nolimits^{þ}$. In particular, we have $U^{þ}(p_0) \leq t$, as desired. $\square$

We can also obtain a converse of 6.6 without too much difficulty.

Proposition 6.7. *Assume that K is rosy, and $U^{þ}$-rank is defined and finite-valued for every type in K. Then $\mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\vert$}}\nolimits^{þ}$ has Local character over K, and K has small algebraicity.*

Proof. To see that $\mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\vert$}}\nolimits^{þ}$ has Local character, it suffices to define $f_{\text{loc}} : \omega \to \omega$ via

$$f_{\text{loc}}(n) = 2 \cdot \max \left\{ U^{þ}(p(\bar{x})) : p(\bar{x}) \in S_n^{\text{qf}}(\emptyset) \right\}$$

It is easily verified that f_{loc} works as advertised. For small algebraicity, we set $d_K = f_{\text{loc}}(1)$. Suppose $\mathcal{M} \in K$, $C \subseteq M$ and $a \in acl(C; \mathcal{M})$. Let $p(x) = qtp(a/C; \mathcal{M})$. By Local character, there is a subset $C_0 \subseteq C$ such that $|C_0| \leq d_K$ and p does not þ-fork over C_0. It follows that $U^{þ}(p \restriction C_0) = U^{þ}(p) = 0$, so $p \restriction C_0$ is algebraic, as required. $\square$

To summarize the results we produced thus far in the chapter, we have the following theorem and a definition:

Theorem 6.8. *Let K be a rosy Fraïssé class of finite structures. The following are equivalent:*

1. *K has small algebraicity*
2. *$\mathop{\smile\!\!\!\!\!\!\mid}^{þ}$ has Local character over K*
3. *$U^{þ}$-rank is defined and finite-valued for every type in K*

If K satisfies any one (hence all) of the above, then we say that K is super-rosy*.*

Besides its existence, the only properties of $U^{þ}$-rank that we will actually require are contained in the following theorem. Its proof is exactly the same as that of the analogous statement in [28], so we forgo recapitulating the demonstration.

Theorem 6.9 (Lascar's equality)**.** *Assume that K is super-rosy. Suppose $\mathcal{M} \in K$, $C \subseteq M$, and $\overline{a}, \overline{b} \in M^{<\omega}$. Then*

$$U^{þ}(\overline{a}, \overline{b}/C; \mathcal{M}) \;=\; U^{þ}(\overline{a}/C; \mathcal{M}) + U^{þ}(\overline{b}/C\overline{a}; \mathcal{M})$$

6.2 þ-Independence is weakest

In this section, we prove a key characterization of þ-independence, namely theorem 6.10. Essentially, the theorem asserts that þ-forking is the most stringent possible interpretation of phrases of the form "... depends on..."; equivalently, $\mathop{\smile\!\!\!\!\!\!\mid}^{þ}$ is the weakest possible interpretation of the phrase "... is independent from..." From this, we can easily derive a characterization of rosy classes that will be fundamental to our analysis of efficient computability of the model-building problem.

Theorem 6.10 (þ-Independence is weakest)**.** *Assume that T is the complete k-variable theory of a finite ρ-structure, which is capped and has infinitely many finite models up to isomorphism, and let $K = \boldsymbol{fin}[T^G]$. Suppose $\mathop{\smile\!\!\!\!\!\!\mid}^{\circ}$ is a notion of independence in K. Then for any $\mathcal{M} \in K$, $A, B, C \subseteq M$, $A \mathop{\smile\!\!\!\!\!\!\mid}^{\circ}_{C} B$ implies $A \mathop{\smile\!\!\!\!\!\!\mid}^{þ}_{C} B$.*

For the proof of theorem 6.10 (which is not significantly different here from the presentation in [10]), we will require a few technical lemmas and observations.

Lemma 6.11. *Let $\mathcal{M} \in K$, $C \subseteq D \subseteq M$, $\overline{a}, \overline{b} \in M^{<\omega}$, and $p(\overline{x}, \overline{b}) = qtp(\overline{a}/C\overline{b}; \mathcal{M})$. If $p(\overline{x}, \overline{b})$ divides strongly over D, then $qtp(\overline{b}/D\overline{a}; \mathcal{M})$ is algebraic, while $qtp(\overline{b}/D; \mathcal{M})$ is non-algebraic.*

Proof. Obviously, $q(\overline{y}) = qtp(\overline{b}/D; \mathcal{M})$ must be non-algebraic by definition of strong dividing. If $q'(\overline{a}, \overline{y}) = qtp(\overline{b}/D\overline{a}; \mathcal{M})$ is non-algebraic, then for every $r < \omega$, there is a model $\mathcal{N} \in K_{D\overline{a}}$ such that the family

$$\left\{q'(\overline{x}, \overline{b}') : \overline{b}' \in N^{\overline{y}}, \mathcal{N} \models q'(\overline{a}, \overline{b}')\right\}$$

is not r-inconsistent, so that $p(\overline{x}, \overline{b})$ cannot divide strongly over D. □

Lemma 6.12. *Suppose $\mathop{\smile\!\!\!\!\!\!|}^{\circ}$ is a notion of independence in K. Let $\mathcal{M} \in K$, $C \subseteq M$, $\overline{a}, \overline{b} \in M^{<\omega}$. If $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_C \overline{b}$, then $p(\overline{x}, \overline{b}) = qtp(\overline{a}/C\overline{b}; \mathcal{M})$ does not þ-divide over C.*

Proof. Assume $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_C \overline{b}$, and towards a contradiction, suppose $p(\overline{x}, \overline{b})$ þ-divides over C. Without loss of generality, we may assume that there is a subset $C \subseteq D \subseteq M$ such that $p(\overline{x}, \overline{b})$ divides strongly over D. In particular, $qtp(\overline{b}/D; \mathcal{M})$ is non-algebraic and for some $1 < r < \omega$, the family

$$\left\{p(\overline{x}, \overline{b}') : \overline{b}' \models qtp(\overline{b}/D; \mathcal{M})\right\}$$

is setwise unbounded and r-inconsistent over K_D. Since $\mathop{\smile\!\!\!\!\!\!|}^{\circ}$, being a notion of independence, has the Extension property and Partial right-transitivity, we may further assume that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_C D\overline{b}$ and, hence, that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_D D\overline{b}$ and $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_D \overline{b}$. By the preceding lemma, we know that $qtp(\overline{b}/D\overline{a}; \mathcal{M})$ is algebraic – that is,

$$\overline{b} \in acl(D\overline{a}; \mathcal{M}) \setminus alg(D; \mathcal{M})$$

So, as $\mathop{\smile\!\!\!\!\!\!|}^{\circ}$ has the property of Preservation of algebraic dependence II, it follows that $\overline{a} \mathop{\not\smile\!\!\!\!\!\!|}^{\circ}_D D\overline{b}$, a contradiction, and this completes the proof of the lemma. □

With these lemmas in hand, we now proceed to the (happily, very concise) proof of the main theorem.

Proof of theorem 6.10. Let $\mathcal{M} \in K$, $\overline{a}, \overline{b} \in M^{<\omega}$ and $C \subseteq M$. For the contrapositive, suppose $\overline{a} \mathop{\not\smile\!\!\!\!\!\!|}^{þ}_C \overline{b}$. By the definition of þ-forking, we may assume that there is a subset $C \subseteq C_1 \subseteq M$ such that every complete extension of $p(\overline{x}, \overline{b}) = qtp(\overline{a}/C\overline{b}; \mathcal{M})$ to $(C_1; \mathcal{M})$ þ-divides over C_1. If $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_C \overline{b}$, then by the Extension property for $\mathop{\smile\!\!\!\!\!\!|}^{\circ}$, we may assume that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_C C_1\overline{b}$, and by Partial right-transitivity, we find that $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\circ}_{C_1} \overline{b}$. By lemma 6.11, it follows, then, that $qtp(\overline{a}/C_1\overline{b}; \mathcal{M})$ *does not* þ-divide over C_1, a contradiction. Thus, $\overline{a} \mathop{\not\smile\!\!\!\!\!\!|}^{\circ}_C \overline{b}$, as desired. □

Corollary 6.13. *Let K be a Fraïssé class of finite structures. The following are equivalent:*

1. *K admits a notion of independence with Local character.*
2. *K is rosy with small algebraicity.*
3. *K is super-rosy.*

Proof. The equivalence of 2 and 3 has already been proven, and 3⇒1 is immediate. For 1⇒2, by proposition 6.1, it suffices to show that if $\mathop{\smile\hspace{-0.9em}\mid}^{\circ}$ is a notion of independence in K with Local character – via $f^{\circ}_{\mathrm{loc}} : \omega \to \omega$ – then $\mathop{\smile\hspace{-0.9em}\mid}^{\text{þ}}$ has Local character as well. Given $\mathcal{M} \in K$, $A, B \subseteq M$, let $B_0 \subseteq B$ such that $|B_0| \leq f^{\circ}_{\mathrm{loc}}(|A|)$ and $A \mathop{\smile\hspace{-0.9em}\mid}^{\circ}_{B_0} B$. By theorem 6.10, we have then $A \mathop{\smile\hspace{-0.9em}\mid}^{\text{þ}}_{B_0} B$, as desired. □

6.2.1 Rosiness (possibly) without super-rosiness

As we've noted before in this chapter, Local character does not seem to be intrinsic to the concept of rosiness for Fraïssé classes of finite structures, and consequently, it would probably be an act of bad conscience to end our discussion with corollary 6.13. Working exclusively with finite sets, however, 6.13 does seem to be the end of what can be recovered, so at this point, it is natural to examine notions of independence *in the classical sense.*

Thus, we define an *independence relation in the classical sense* – or, more briefly, a classical independence relation – in the same manner as we've been using, except that we eliminate our restrictions to finite sets, and we *do not require the Base-monotonicity property.* For clarity, the concept we have (and will be) particularly interested in will be called, temporarily, a *finitary notion of independence* or a *finitary independence relation.* (Note that a finitary notion of independence is only defined for triples of subsets of finite structures.) We now show how to lift a finitary notion of independence to a classical one. Suppose $\mathop{\smile\hspace{-0.9em}\mid}^{\circ}$ is a finitary notion of independence on the Fraïssé class K, and let $T^c = Th(\mathfrak{M})$, where $\mathfrak{M}$ is the countable Fraïssé limit of K. Note that T^c is countably categorical and eliminates quantifiers. Let $\mathfrak{M}^*$ be the monster model of T^c. For $A, B, C \subset ||\mathfrak{M}^*||$, we assert $A \widehat{\mathop{\smile\hspace{-0.9em}\mid}^{\circ}}_C B$ just in case the following condition obtains:

> For all $\overline{a} \in A^{<\omega}$, there is a $C_0 = C_0(\overline{a})$ such that for all $\overline{b} \in B^{<\omega}$, if $C_0 \subseteq D \subseteq_{\mathbf{fin}} C$, then $\overline{a} \mathop{\smile\hspace{-0.9em}\mid}^{\circ}_D \overline{b}$.

Proposition 6.14. *If $\mathop{\smile\hspace{-0.9em}\mid}^{\circ}$ is a finitary notion of independence, then $\widehat{\mathop{\smile\hspace{-0.9em}\mid}^{\circ}}$ is a classical notion of independence with the Existence property.*

Proof. Invariance, Monotonicity, Base-monotonicity, Partial right-transitivity, and Preservation of algebraic dependence I-III are all easy to verify. Thus, it remains only to demonstrate that $\widehat{\mathop{\smile\hspace{-0.9em}\mid}^{\circ}}$ has the Extension property.

Suppose $\overline{a} \widehat{\mathop{\smile\hspace{-0.9em}\mid}^{\circ}}_C B$, where $\overline{a}$ is a finite tuple, and let $B \subseteq B_1$. For $\overline{b}_1 \in B_1^{<\omega}$ and a partial type $\pi(\overline{x}, \overline{b}_1)$ over $B_1 C$ consistent with $qtp(\overline{a}/BC)$, we say that $\pi(\overline{x}, \overline{b}_1)$ is a *bad type* if there is a finite subset $C_0 \subseteq D \subseteq_{\mathbf{fin}} C$ such that if $\overline{a}' \vDash \pi(\overline{x}, \overline{b}_1)$, then $\overline{a}' \mathop{\smile\hspace{-0.9em}\mid}^{\circ}_D \overline{b}_1$. By the invariance property of $\mathop{\smile\hspace{-0.9em}\mid}^{\circ}$ and compactness, every bad type has a finite subtype which is also bad. Using the Extension property of $\mathop{\smile\hspace{-0.9em}\mid}^{\circ}$, a routine compactness argument shows that for some $p(\overline{x}) \supseteq qtp(\overline{a}/BC)$ over $B_1 C$, if $\overline{a}' \vDash p$, then $\overline{a}' \widehat{\mathop{\smile\hspace{-0.9em}\mid}^{\circ}}_C B_1$. □

Proposition 6.15. *If $\unicode{x2ADD}^{\circ}$ is a finitary true independence relation, then $\widehat{\unicode{x2ADD}}^{\circ}$ is a classical true independence relation.*

Proof. We need only verify that $\widehat{\unicode{x2ADD}}^{\circ}$ is symmetric and transitive. The former is obvious from the definition of $\widehat{\unicode{x2ADD}}^{\circ}$ under the assumption that $\unicode{x2ADD}^{\circ}$ is symmetric. For transitivity, suppose $\overline{a}\widehat{\unicode{x2ADD}}^{\circ}_{C}B_1$ and $\overline{a}\widehat{\unicode{x2ADD}}^{\circ}_{CB_1}B_2$; we must show that $\overline{a}\widehat{\unicode{x2ADD}}^{\circ}_{C}B_1B_2$. First, we make two simple observations – the first easily follows from the transitivity of $\unicode{x2ADD}^{\circ}$, and the second immediately from the symmetry of $\unicode{x2ADD}^{\circ}$.

Observation. *If $\overline{b}_1, \overline{b}_2$ are finite and $\overline{a}\widehat{\unicode{x2ADD}}^{\circ}_{C}\overline{b}_1 \wedge \overline{a}\widehat{\unicode{x2ADD}}^{\circ}_{C\bar{b}_1}\overline{b}_2$, then $\overline{a}\widehat{\unicode{x2ADD}}^{\circ}_{C}\overline{b}_1\overline{b}_2$.*

Observation. *For arbitrary sets A, B, C, $A\widehat{\unicode{x2ADD}}^{\circ}_{C}B$ if and only if $\overline{a}\widehat{\unicode{x2ADD}}^{\circ}_{C}\overline{b}$ for all $\overline{a} \in A^{<\omega}$ and $\overline{b} \in B^{<\omega}$.*

We now complete the proof of the fact that $\widehat{\unicode{x2ADD}}^{\circ}$ is fully transitive. For the contrapositive, note that if $\overline{a}\widehat{\not\unicode{x2ADD}}^{\circ}_{C}B_1B_2$, then $B_1B_2\widehat{\not\unicode{x2ADD}}^{\circ}_{C}\overline{a}$, so that $\overline{b}1\overline{b}_2\widehat{\not\unicode{x2ADD}}^{\circ}_{C}\overline{a}$ for some $\overline{b}_1 \in B_1^{<\omega}$ and $\overline{b}_2 \in B_2^{<\omega}$; thus, by the first observation, $\overline{a}\widehat{\not\unicode{x2ADD}}^{\circ}_{C}\overline{b}_1$ or $\overline{a}\widehat{\not\unicode{x2ADD}}^{\circ}_{C\bar{b}_1}\overline{b}_2$, and this completes the proof. $\square$

Now, in [10], those authors prove the following theorems, the first being the direct analog of our theorem 6.10. The proof of the second theorem seems to depend strongly on work of the first author's thesis, through which forking and þ-forking are characterized in terms of finite satisfiability of extensions of types. "Finite-satisfiability" is not a useful condition for the analysis of finitary notions of independence, and of course, this was the reason for our examination of classical notions of independence in the first place.

Theorem 6.16 (þ-Independence is weakest, classical version: 3.3 of [10])**.** *Suppose $\unicode{x2ADD}^{*}$ is a classical notion of independence for a theory T (on real subsets of models). Then, for all real subsets A, B, C, $A\unicode{x2ADD}^{*}_{C}B$ implies $A\unicode{x2ADD}^{þ}_{C}B$.*

Theorem 6.17 (3.8 of [10])**.** *Suppose $\unicode{x2ADD}^{*}$ is a classical notion of independence for a theory T (on real subsets of models).*

1. *If $\unicode{x2ADD}^{*}$ has Local character (in the classical sense), then T is real-rosy and $\unicode{x2ADD}^{þ}$ is a classical true independence relation (for real subsets of models).*

2. *If $\unicode{x2ADD}^{*}$ is symmetric and fully transitive, then T is real-rosy.*

As a corollary, we can partially recover the characterization of rosiness from [10]. The recovery is partial, for as we compose this dissertation, it does not seem easy to show that the classical Local character of the canonical limit theory T^c pushes down to Local character in the sense we defined above

Lemma 6.18. *Let K be a Fraïssé class of finite structures, and let $T^c = Th(\mathfrak{M})$, where $\mathfrak{M}$ is the countable Fraïssé limit of K. If T^c is (real) rosy, then K is rosy.*

Proof. Assuming T^c is rosy, we know that $\text{þ}(\overline{x} = \overline{x}, \Phi, \Delta, r) < \omega$ is always defined, which, of course, means that K is rosy. □

Corollary 6.19. *Let K be a Fraïssé class of finite structures. Let $\mathop{\smile\!\!\!\!\!\!|}^{\circ}$ be a finitary notion of independence in K. If $\mathop{\smile\!\!\!\!\!\!|}^{\circ}$ is finitary true independence relation, then K is rosy.*

Proof. If $\mathop{\smile\!\!\!\!\!\!|}^{\circ}$ is symmetric and transitive, then its canonical lift $\widehat{\mathop{\smile\!\!\!\!\!\!|}^{\circ}}$ is symmetric and transitive. Thus, T^c, the canonical limit theory, is real-rosy, and it follows from the lemma, that K is rosy. □

6.3 Examples and non-examples

In this section, we will present a few slightly interesting examples – two of rosy finite-variable theories and one non-rosy theory. This list is, obviously, not at all exhaustive, and to remove any suspense, we state now the kinds of examples we have not been able to find thus far:

1. Are there rosy classes which are not super-rosy?

   In the context of capped finite-variable theories, this amounts to confirming or refuting the assertion, "all capped finite-variable theories have small algebraicity." To us, it seems most likely that this assertion is true – and indeed we have not found any natural theories that do not have small algebraicity – but a proof, so far, eludes us. In the more general context of an arbitrary Fraïssé class, we have no strong opinions in either direction.

2. Are there rosy classes which are not simple?

   In the full first-order world, this question is not difficult at all, for every o-minimal theory is rosy and un-simple. In k-variable logic with $k \geq 3$, an ordered finite graph, say, is determined up to isomorphism by its k-variable theory. Thus, we don't have recourse to this easy example. In future work, we expect to examine the k-variable theory of the triangle-free random graph and/or large bipartite digraphs.

Before we begin with specific examples, we present a proposition which will be essential in the later exposition. Let ρ be a finite relational signature, and assume that $k \geq ari(R)$ for all $R \in \rho$ (and in any case, $k \geq 2$). Suppose $l < k$, and let $\theta(x_1, ..., x_l, y)$ be a complete quantifier-free $l+1$-type of ρ. Then, let $\theta_0(x_1, ..., x_l)$ unique quantifier-free l-type such that $\theta \vDash \theta_0$, which call the type induced on $x_1, ..., x_l$ by θ. For brevity in the sequel, we write $\theta_{\restriction(x_1,...,x_l)}$ for θ_0. The *extension axiom over* θ is the sentence,

$$E_\theta : \ \forall x_1, ..., x_l \, (\theta_0(x_1, ..., x_l) \to \exists y \, \theta(\overline{x}, y))$$

If T is a k-variable theory, we say that θ is *compatible* with T if $T \cup \{\exists \overline{x}, y\, \theta\}$ has a (not necessarily finite) model. We lift the following theorem directly from [12].

Proposition 6.20 (Extension axioms)**.** *Let T be a consistent k-variable theory. The following are equivalent*

1. *T is complete and eliminates quantifiers.*
2. *$T \vDash E_\theta$ for every quantifier-free $l+1$-type, $l < k$, $\theta(\overline{x}, y)$ which is compatible with T.*

6.3.1 "Large" vector spaces over finite fields

For this section, we fix a finite $\mathscr{F}$. The signature for $\mathscr{F}$-vector spaces is the following:

$$\rho_{vs} = \left\{P_0^{(1)}, R_+^{(3)}\right\} \cup \left\{R_\alpha^{(2)} : \alpha \in \mathscr{F}\right\}$$

If $\mathcal{M}$ is an $\mathscr{F}$-vector space, we can easily interpret it as a ρ_{vs}-structure:

- $P_0^{\mathcal{M}} = \{0_{\mathcal{A}}\}$
- $R_+^{\mathcal{M}} = \{(a, b, c) \in M^3 : a + b = c\}$
- $R_\alpha^{\mathcal{M}} = \{(b, \alpha b) : b \in M\}$

We saw in chapter 2, that if $\mathcal{A}$ is of sufficiently high dimension over $\mathscr{F}$, then $T = Th^k(\mathcal{M}_0)$ is a capped theory with infinitely many models up to isomorphism. Computing algebraic closures in the sense of T, it is relatively straightforward to show that the following is an independence relation in $K = \mathbf{fin}[T^G]$: Suppose $\mathcal{M} \in K$, $A, B, C \subseteq M$; then

$$A \mathop{\smile\!\!\!\!\!\!|}^{\circ}_C B \Leftrightarrow \langle AC \rangle \cap \langle BC \rangle = \langle C \rangle \text{ in } \mathcal{N}^{\mathbf{mod}} \text{ for some } \mathcal{N} \in K_{ABC}$$

where $\langle \cdot \rangle$ denotes linear span. It follows that K is a rosy class. It's also easy to show that K has small algebraicity, so in fact, K is super-rosy.

6.3.2 The random graph in k-variables

As usual, the signature graphs ρ_{gr} has a single relation symbol $R^{(2)}$. For $k \geq 2$, we define $\mathcal{RG}_k$ to be the theory consisting of the symmetry axiom, $\forall xy(R(x, y) \to R(y, x))$, and the following extension axioms,

$$E_{l,S}: \ \forall x_1, ..., x_l \left(\bigwedge_{i<j} x_i \neq x_j \to \exists y \left(\bigwedge_{i \in S} R(x_i, y) \wedge \bigwedge_{i \in [l] \setminus S} \neg R(x_i, y) \right) \right)$$

for each $l < k$ and $S \subseteq [l]$. By proposition 6.20, $\mathcal{RG}_k$ is a complete k-variable with elimination of quantifiers. Moreover, $\mathcal{RG}_k^G$ amounts to specifying the complete quantifier-free k-types, so there is *really* nothing to distinguish the two.

As for the full first-order theory of the random graph, it can be show that if $\mathcal{M} \models \mathcal{RG}_k$, $A \subseteq M$ and $b \in M$ is algebraic over A, then $b \in A$. Hence, $\mathcal{RG}_k$ has small algebraicity, and it follows, without further analysis, that $\mathop{\smile\!\!\!\!\!\!\,|}^{\text{þ}}$ has Local character, so that $\mathcal{RG}_k$ is super-rosy.

6.3.3 Random pseudo-scales – an analog of parametrized equivalence relations

One of the most natural examples of a non-rosy theory in the classical sense is sometimes called the theory of parametrized equivalence relations, which for clarity in the rest of this section, we now present. Let $\mathcal{M}$ be the structure with two sorts $S_p = \omega$ and $S_f = {}^\omega\omega$ (points and functions), and let $E^\mathcal{M} \subseteq {}^\omega\omega \times {}^\omega\omega \times \omega$ be the family of equivalence relations given by

$$E^\mathcal{M} = \{(g_1, g_2, a) : g_1(a) = g_2(a)\}$$

Then, it is routine to verify that $\text{þ}(S_f(x), E(x, x', y), S_p(y), 2) = \infty$. The objective of this section is to define a 3-variable theory which, essentially, simulates this classical example, and which is still capped. In order provably to ensure that the latter condition is satisfied, we must introduce a certain amount of complexity into the construction. Fortunately, much of the analysis has already been carried out in [12], and our presentation is a minor perturbation of that, so we omit the proofs. The terms *scale* and *pseudo-scale* come from that paper as well, although we do not concern ourselves explicitly enforcing bijectivity anywhere.

The signature for scales and pseudo-scales, like that for parametrized equivalence relations, is to be understood as encoding a family of functions from one sort to another. In order to economize on variables, we have an additional sort consisting of the "names" of the functions, and we identify functions with their graphs. Thus, the signature of scales and pseudo-scales is,

$$\rho_{sc} = \left\{ S_1^{(1)}, S_2^{(1)}, Q^{(1)}, F^{(1)}, P_1^{(2)}, P_2^{(2)}, E_1^{(2)}, E_2^{(2)}, R_\in^{(2)} \right\}$$

and the "natural" interpretation of these in a model $\mathcal{M}$ is as follows:

1. $S_1^\mathcal{M}$, $S_2^\mathcal{M}$ are disjoint sets, and $Q^\mathcal{M} = S_1^\mathcal{M} \times S_2^\mathcal{M}$
2. $P_1^\mathcal{M} = \{((a_1, a_2), a_1) : (a_1, a_2) \in Q^\mathcal{M}\} \subseteq Q^\mathcal{M} \times S_1^\mathcal{M}$
   $P_2^\mathcal{M} = \{((a_1, a_2), a_2) : (a_1, a_2) \in Q^\mathcal{M}\} \subseteq Q^\mathcal{M} \times S_2^\mathcal{M}$
3. $E_1^\mathcal{M} = \{((a_1, a_2), (b_1, b_2)) \in Q^\mathcal{M} \times Q^\mathcal{M} : a_1 = b_1\}$
   $E_2^\mathcal{M} = \{((a_1, a_2), (b_1, b_2)) \in Q^\mathcal{M} \times Q^\mathcal{M} : a_2 = b_2\}$

4. $F^{\mathcal{M}}$ is a non-empty set of functions $S_1^{\mathcal{M}} \to S_2^{\mathcal{M}}$
5. $R_{\in}^{\mathcal{M}} = \left\{((a,b),f) \in Q^{\mathcal{M}} \times F^{\mathcal{M}} : f(a) = b\right\}$

Such a structure is called a *natural scale*, and any structure isomorphic to some natural scale is called a *scale*.

Lemma 6.21. *The class of scales is axiomatizable by a single sentence of $L(\rho_{sc})^3$.*

For parameters $0 < n, l, \omega$, we define a probability space of natural *pseudo-scales*, $W(n,l)$, as follows. For each $\mathcal{M} \in W(n,l)$,

- $S_1^{\mathcal{M}}$ and $S_2^{\mathcal{M}}$ are disjoint copies of $[n]$, and $Q^{\mathcal{M}} = S_1^{\mathcal{M}} \times S_2^{\mathcal{M}}$
- $P_1^{\mathcal{M}}$, $P_2^{\mathcal{M}}$, $E_1^{\mathcal{M}}$, $E_2^{\mathcal{M}}$ have their natural interpretations, and $F^{\mathcal{M}}$ will have its natural interpretation once $F^{\mathcal{M}}$ is specified.
- $F^{\mathcal{M}} = \left\{f_1^{\mathcal{M}}, ..., f_l^{\mathcal{M}}\right\}$ is a *multiset* of function *names*, and for each $i \in [l]$, we choose $f_i : [n] \to [n]$ independently and uniformly at random.

We observe, then, that

$$\begin{aligned}\mathbb{P}\left\{\mathcal{M} \in W(n,l) : \bigvee_{i<j} f_i^{\mathcal{M}} = f_j^{\mathcal{M}}\right\} &\leq l^2 \cdot \mathbb{P}\left\{(f,g) \in ({}^{[n]}[n])^2 : f = g\right\} \\ &\leq l^2 \cdot \frac{n^n}{(n^n)^2} \\ &= \frac{l^2}{n^n}\end{aligned}$$

Hence, take $l(n) = n^5$, say, we have

$$\lim_{n\to\infty} \mathbb{P}\left\{\mathcal{M} \in W(n,n^5) : \bigvee_{i<j} f_i^{\mathcal{M}} = f_j^{\mathcal{M}}\right\} = 0$$

Equivalently,

$$\lim_{n\to\infty} \mathbb{P}\left\{\mathcal{M} \in W(n,n^5) : \mathcal{M} \text{ is a natural scale}\right\} = 1$$

Now, for a pseudo-scale $\mathcal{M}$, let $\varphi^*(x_1,x_2)$ be the following L^3-formula.

$$(F(x_1) \wedge F(x_2) \wedge x_1 \neq x_2) \wedge \exists x_3(Q(x_3) \wedge R_{\in}(x_3,x_1) \wedge R_{\in}(x_3,x_2))$$

Thus, in a pseudo-scale, φ^* defines the set,

$$D^{\mathcal{M}} = \left\{(f,g) \in F^{\mathcal{M}} \times F^{\mathcal{M}} : f \neq g \wedge \exists x(f(x) = g(x))\right\}$$

and we expand the signature ρ_{sc} to $\rho_{sc}^* = \rho_{sc} \cup \{D^{(2)}\}$ and adjoin an additional axiom: $\varphi_D = \forall x_1 x_2(D(x_1,x_2) \leftrightarrow \varphi^*(x_1,x_2))$. Now, the following proposition is found in [12]:

Proposition 6.22 (Lemma 14 of [12]). *There is an L^3-theory T of the signature ρ_{sc} such that the following hold:*

1. *T is complete for L^3, and the L^3-theory of the class of expansions*

$$\{(\mathcal{M}, D^{\mathcal{M}}) : \mathcal{M} \vDash T\}$$

   is complete and eliminates quantifiers.

2. *There is a number $n_0 < \omega$ such that if $n \geq n_0$, then there is a natural scale $\mathcal{M}$ such that $S_1^{\mathcal{M}} = S_2^{\mathcal{M}} = [n]$, $|F^{\mathcal{M}}| = n^5$ and $\mathcal{M} \vDash T$.*

We take it as obvious that T has small algebraicity, and the following proposition is proved using almost identical methods:

Proposition 6.23. *T is capped. More precisely, suppose $\mathcal{A}$ is a finite model of $T_\forall^G$ with $n_1 = |S_1^{\mathcal{A}^{mod}}|$. Then, for sufficiently large $n_1 \leq n < \omega$,*

$$\frac{\mathbb{P}\{\mathcal{M} \in W(n, n^5) : \mathcal{A} \leq \mathcal{M} \wedge \mathcal{M} \vDash T\}}{\mathbb{P}\{\mathcal{M} \in W(n, n^5) : \mathcal{A} \leq \mathcal{M}\}} > 0$$

We note, however, that there is no guarantee that the capping model of an induced diagram $\mathcal{A}$ is of bounded size. In fact, the argument proceeds by taking n so much larger than n_1 that replacing $W(n, n^5)$ with another probability space $W_{\mathcal{A}}(n, n^5)$ of pseudo-scales guaranteed to preserve $\mathcal{A}$ is inconsequential. In any case, the following observation, demonstrating that T is not rosy, is fairly easy to see:

Observation. *If $\varphi(x_1, x_2, x_3) = R_\in(x_3, x_1) \wedge R_\in(x_3, x_2)$ and $\delta(x_2, x_3) = S_1(x_2) \wedge P(x_3)$, then $\text{þ}(F(x_1), \varphi(x_1, x_2), \delta(x_2, x_3), 2) = \infty$.*

Chapter 7

Coordinatization and efficient model-building

Throughout this chapter, we take K to be a super-rosy class of finite ρ-structures which has amalgamation over sets and which eliminates imaginaries. The class K^{eq}, where $K = \mathbf{fin}[T^G]$ and T is a finite-variable theory fits this description. We will make rather heavy use in this chapter of the $U^{\mathrm{þ}}$-rank (defined below), especially of the Lascar equality (see chapter 6). Under the assumption of super-rosiness, we show that K admits a polynomial-time solution of its model-building problem.

For the sake of brevity, we will not address the issue of resolving algebraic types. However, by the super-rosiness assumption – i.e. small algebraicity – it should be clear that such resolution is tractable in relational polynomial-time, as first-order logic, whence `invent`-programs, can "count" up to an *a priori* fixed number, and this can be hard-coded into the program. (If we wished to recover the program itself by algorithmic means, we could not gloss this point.) Thus, the algorithm in our presentation amounts to proving a weak model-theoretic coordinatization theorem and using this to keep track of the flow of information – more specifically, keep track of the need to invent new elements – in the model-building process.

7.1 Coordinatization machinery

In this section, we define a notion of coordinatization in the model-theoretic sense. Our result is significantly weaker than those recovered in [6, 7], [9], or even [21], as the "coordinatizing object" we obtain is not a tree, nor even acyclic in fact. We note, for the sake of interest, that it is possible to obtain a somewhat more typical coordinatization result, but this is result is not terribly useful algorithmically.

A further note on the content of this section is in order. The material in subsection 7.1.1 was discovered independently by the present author in the late summer of 2009, but it has since been published in [23], which appeared on `arXiv.org` in October 2009. The presentation given below is a compromise between the notation of our derivation and that in [23], and part 3 of the definition of a coordinate was not part of our independent result.

7.1.1 Finding coordinates of tuples

Let $\mathcal{M} \in K$, $D \subseteq M$ and $\overline{a}, \overline{b}, \overline{e} \in M^{<\omega}$, and assume $U^{\text{þ}}(\overline{a}/D) = \alpha + 1$. We say that the pair $(\overline{b}, \overline{e})$ is a *coordinate of* $\overline{a}$ *over* D if the following conditions are satisfied:

1. $\overline{a} \downarrow^{\text{þ}}_{D} \overline{e}$ and $\overline{a} \downarrow^{\text{þ}}_{D\overline{b}} \overline{e}$;
2. $U^{\text{þ}}(\overline{a}/D\overline{b}) = \alpha$ and $qtp(\overline{a}, \overline{b}/D; \mathcal{M})$ divides strongly over $D\overline{e}$ in $qtp(\overline{b}/D\overline{e})$ (in particular, $qtp(\overline{b}/D\overline{a}\overline{e}; \mathcal{M})$ is algebraic);
3. $U^{\text{þ}}(\overline{b}/D\overline{e}) = 1$.

Lemma 7.1. *Let $\mathcal{M}_0 \in K$, $C \subseteq M_0$, $\overline{a} \in M_0^{<\omega}$ and $p(\overline{x}) = qtp(\overline{a}/C; \mathcal{M}_0)$, and suppose $U^{þ}(p) = \alpha + 1$. Then there are $\mathcal{M} \in K_{C\overline{a}}$, $\overline{b} \in M^{<\omega}$, and an extension $q(\overline{x}, \overline{b})$ of $p(\overline{x})$ to $C\overline{b}$ such that $U(q(\overline{x}, \overline{b})) = \alpha$ and $q(\overline{x}, \overline{b})$* þ-divides *over C.*

Proof. By the definition of $U^{\text{þ}}$-rank, let $\mathcal{M} \in K_C$ and $C \subsetneq C_1 \subseteq M$, and let $p_1(\overline{x})$ be a complete extension of $p(\overline{x})$ to C_1 such that $U^{\text{þ}}(p_1) = \alpha$ and p_1 þ-forks over C. Since p_1 þ-forks over C, there are $\overline{b}_1, ..., \overline{b}_s \in C_1^{<\omega}$ and formulas $\varphi_1(\overline{x}, \overline{y}_1)$, ..., $\varphi_s(\overline{x}, \overline{y}_s)$ such that

$$K_C \models p_1(\overline{x}) \to \bigvee_{i=1}^{s} \varphi_i(\overline{x}, \overline{b}_i)$$

and each $\varphi_i(\overline{x}, \overline{b}_i)$ þ-divides over C. As p_1 is a complete type, we may assume, without loss of generality, that $\varphi_1(\overline{x}, \overline{b}_1) \in p_1$, and we set $q(\overline{x}, \overline{b}_1) = p_1 \restriction C\overline{b}_1$. Then $q(\overline{x}, \overline{b}_1)$ þ-divides over C, and

$$\alpha = U^{\text{þ}}(p_1) \leq U^{\text{þ}}(q(\overline{x}, \overline{b}_1)) < U^{\text{þ}}(p) = \alpha + 1$$

so that $U^{\text{þ}}(q(\overline{x}, \overline{b}_1)) = \alpha$, as desired. □

Lemma 7.2. *Let $\mathcal{M}_0 \in K$, $C \subseteq M_0$ and $\overline{a}, \overline{b} \in M_0^{<\omega}$, and suppose $qtp(\overline{a}/C\overline{b})$ þ-divides over C. Then there are $\mathcal{M} \in K_{C\overline{a}\overline{b}}$ and $\overline{e} \in M^{<\omega}$ such that*

1. *$\overline{a} \downarrow^{þ}_{C\overline{b}} \overline{e}$*
2. *$qtp(\overline{a}/C\overline{b}\overline{e}; \mathcal{M})$ divides strongly over $C\overline{e}$ in $qtp(\overline{b}/C\overline{e}; \mathcal{M})$ (so $qtp(\overline{b}/C\overline{a}\overline{e}; \mathcal{M})$ is algebraic)*
3. *$U^{þ}(\overline{a}/C\overline{b}\overline{e}; \mathcal{M}) < U^{þ}(\overline{a}/C\overline{e}; \mathcal{M})$*

Proof. Let $\mathcal{N} \in K_{C\overline{b}}$ and $\overline{d} \in N^{<\omega}$ such that $qtp(\overline{a}/C\overline{b};\mathcal{M})$ divides strongly over $C\overline{d}$ in $qtp(\overline{b}/C\overline{d};\mathcal{N})$. Observe that if $qtp(\overline{b}/C\overline{d})$ is algebraic, then it cannot possibly induce strong dividing, so $qtp(\overline{b}/C\overline{d})$ must be non-algebraic. Without loss of generality, we may assume that $\mathcal{N} \in K_{C\overline{a}\overline{b}}$. Then $\overline{d}\mathop{\smile\!\!\!\!\!\!\mid}^{\text{þ}}_{C\overline{b}} C\overline{b}$ by Existence, so by Extension, Symmetry and Monotonicity, we obtain $\mathcal{M} \in K_{C\overline{a}\overline{b}}$ and $\overline{e} \in M^{<\omega}$ such that $qtp(\overline{e}/C\overline{b};\mathcal{M}) = qtp(\overline{d}/C\overline{b};\mathcal{N})$ and $\overline{a}\mathop{\smile\!\!\!\!\!\!\mid}^{\text{þ}}_{C\overline{b}}\overline{e}$. Clearly, $qtp(\overline{a}/C\overline{b}\overline{e};\mathcal{M})$ divides strongly over $C\overline{e}$ in $qtp(\overline{b}/C\overline{e};\mathcal{M})$, and in particular, $qtp(\overline{a}/C\overline{b}\overline{e};\mathcal{M})$ þ-forks over $C\overline{e}$, which necessitates $U^{\text{þ}}(\overline{a}/C\overline{b}\overline{e};\mathcal{M}) < U^{\text{þ}}(\overline{a}/C\overline{e};\mathcal{M})$. □

Proposition 7.3 (Coordinate-finding lemma)**.** *Let $\mathcal{M}_0 \in K$, $C \subseteq M_0$ and $\overline{a} \in M_0^{<\omega}$, and suppose $U^{þ}(\overline{a}/C) > 1$. Then there are $\mathcal{M} \in K_{C\overline{a}}$ and $\overline{b},\overline{e} \in M^{<\omega}$ such that $(\overline{b},\overline{e})$ is a coordinate of $\overline{a}$ over C.*

Proof. Let $p(\overline{x}) = qtp(\overline{a}/C;\mathcal{M}_0)$, so that $U^{\text{þ}}(p) = \alpha + 1 > 1$. By lemma 7.1, we select $\mathcal{M}_1 \in K_C$, $\overline{b} \in M_1^{<\omega}$ and a complete extension $p_1(\overline{x},\overline{b})$ of $p(\overline{x})$ such that $U^{\text{þ}}(p_1(\overline{x},\overline{b})) = \alpha$ and $p_1(\overline{x},\overline{b})$ þ-divides over C. By 7.2, there are $\mathcal{M}_2 \in K_{C\overline{a}\overline{b}}$ and $\overline{e} \in M_2^{<\omega}$ such that $\overline{a}\mathop{\smile\!\!\!\!\!\!\mid}^{\text{þ}}_{C\overline{b}}\overline{e}$, $qtp(\overline{a}/C\overline{e}\overline{b};\mathcal{M}_2)$ divides strongly over $C\overline{e}$ in $qtp(\overline{b}/C\overline{e};\mathcal{M}_2)$ (so $\overline{b} \in acl(C\overline{ae})$) and $U^{\text{þ}}(\overline{a}/C\overline{b}\overline{e};\mathcal{M}_2) < U^{\text{þ}}(\overline{a}/C\overline{e};\mathcal{M}_2)$. Now,

$$\begin{aligned}\alpha &= U^{\text{þ}}(\overline{a}/C\overline{b};\mathcal{M}_1)\\ &< U^{\text{þ}}(\overline{a}/C\overline{b}\overline{e};\mathcal{M}_2)\\ &\leq U^{\text{þ}}(\overline{a}/C\overline{e};\mathcal{M}_2)\\ &= U^{\text{þ}}(\overline{a}/C;\mathcal{M}_0)\\ &= \alpha+1\end{aligned}$$

so $U^{\text{þ}}(\overline{a}/C\overline{e};\mathcal{M}_2) = \alpha + 1 = U^{\text{þ}}(\overline{a}/C;\mathcal{M}_0)$; it follows that $\overline{a}\mathop{\smile\!\!\!\!\!\!\mid}^{\text{þ}}_{C}\overline{e}$ from the definition of $U^{\text{þ}}$-rank. It remains only to show that $U^{\text{þ}}(\overline{b}/C\overline{e};\mathcal{M}_2) = 1$.

By the Lascar inequality and since $U^{\text{þ}}(\overline{b}/C\overline{ae};\mathcal{M}_2) = 0$ (because $qtp(\overline{b}/C\overline{ae};\mathcal{M}_2)$ is algebraic), then,

$$\begin{aligned}U^{\text{þ}}(\overline{a}\overline{b}/C\overline{e};\mathcal{M}_2) &= U^{\text{þ}}(\overline{a}/C\overline{e};\mathcal{M}_2) + U^{\text{þ}}(\overline{b}/C\overline{ae};\mathcal{M}_2)\\ &= (\alpha+1) + U^{\text{þ}}(\overline{b}/C\overline{ae};\mathcal{M}_2)\\ &= \alpha+1\end{aligned}$$

Again, applying the Lascar inequality,

$$\begin{aligned}U^{\text{þ}}(\overline{b}/C\overline{ae};\mathcal{M}_2) + U^{\text{þ}}(\overline{a}/C\overline{e};\mathcal{M}_2) &= U^{\text{þ}}(\overline{a}\overline{b}/C\overline{e};\mathcal{M}_2)\\ &= U^{\text{þ}}(\overline{a}/C\overline{e};\mathcal{M}_2) + U^{\text{þ}}(\overline{b}/C\overline{ae};\mathcal{M}_2)\end{aligned}$$

so that

$$\alpha + U^{\text{þ}}(\overline{b}/C\overline{e};\mathcal{M}_2) \leq \alpha + 1 \leq \alpha + U^{\text{þ}}(\overline{b}/C\overline{e};\mathcal{M}_2)$$

Thus, $U^{\text{þ}}(\overline{b}/C\overline{e};\mathcal{M}_2) = 1$, as desired. □

7.2 Self-coordinatized systems

Our notion of coordinatization in this section, specifically the idea of a self-coordinatized set, is based on the very similar notion in [21], where it is deployed to understand $\aleph_0$-categorical simple theories with 1-based trivial non-forking independence. (Here, since we surrender the requirement that the coordinatizing object is tree, we can do without the 1-basedness hypothesis, and since we do not care whether algebraic closure is a pregeometry on the self-coordinatized set, we can also do without trivial independence.) Intuitively speaking, perhaps the best way to think about our self-coordinatized systems is as a systemic way of obtaining models through iterating the algebraic-closure operation.

In this section particularly, it will be most convenient to work in a (countably) infinite Fraïsse limit $\mathfrak{M}$ of the class K. Thus, if $\overline{a} \in ||\mathfrak{M}||^{<\omega}$ and $C \subseteq ||\mathfrak{M}||$, then we have

$$U^\flat(\overline{a}) = U^\flat(\overline{a}; \mathfrak{M}) = U^\flat(qtp(\overline{a}; \mathfrak{M}))$$

and

$$U^\flat(\overline{a}/C) = U^\flat(\overline{a}/C; \mathfrak{M}) = U^\flat(qtp(\overline{a}/C; \mathfrak{M}))$$

We also write $\overline{a} \equiv^{\mathrm{qf}} \overline{b}$ to mean that $qtp(\overline{a}; \mathfrak{M}) = qtp(\overline{b}; \mathfrak{M})$. We fix $t < \omega$ such that $t \geq ari(R)$ for all $R \in \rho$, and in any case $t \geq 2$. So as to forestall some burdensome notation, we will suppress most explicit references to t in the sequel.

Let $t \leq m < \omega$ and let $A, L \subseteq ||\mathfrak{M}||^{\leq m}$ such that $L \subseteq A \cap ||\mathfrak{M}||^t$. For $\overline{a} \in ||\mathfrak{M}||^{<\omega}$, we write $\overline{a} \sqsubseteq L$ if there are $\overline{c}_1, ..., \overline{c}_n \in L$ such that $rng(\overline{a}) = rng(\overline{c}_1) \cup \cdots \cup rng(\overline{c}_n)$. We say that the pair (A, L) is *$\sqsubseteq$-reflexive* if $\overline{e} \sqsubseteq L$ for every $\overline{e} \in E$, and we say that A is *flush* if for all $\overline{a}, \overline{b} \in ||\mathfrak{M}||^{\leq m}$, if $\overline{a} \in A$ and $\overline{a} \equiv^{\mathrm{qf}} \overline{b}$, then $\overline{b} \in A$.

Now, assume that $A, E \subseteq ||\mathfrak{M}||^{\leq m}$ are flush, $L \subseteq A \cap ||\mathfrak{M}||^t$ is also flush, and $(A \cup E, L)$ is $\sqsubseteq$-reflexive, and consider a partial function

$$\texttt{crd} : A \rightharpoonup (A \times E) \cup \{\star\}$$

Suppose that crd satisfies the following conditions:

1. If $\overline{a} \in dom(\texttt{crd})$ and $\overline{a} \equiv^{\mathrm{qf}} \overline{a}'$, then $\overline{a}' \in dom(\texttt{crd})$.
2. If $\overline{a} \in A$ and $U^\flat(\overline{a}) \leq 1$, then $\texttt{crd}(\overline{a}){\downarrow} = \star$.
3. If $\overline{a} \in dom(\texttt{crd})$ and $U^\flat(\overline{a}) > 1$, then $(\overline{b}, \overline{e}) = \texttt{crd}(\overline{a})$ is a coordinate of $\overline{a}$ over $\emptyset$ such that $rng(\overline{b}) \cap acl(\overline{e}) = \emptyset$ and $rng(\overline{a}) \cap rng(\overline{b}) = \emptyset$.
4. If $\overline{a}, \overline{a}' \in dom(\texttt{crd})$ and $\overline{a} \equiv^{\mathrm{qf}} \overline{a}'$, then $\overline{a}^\frown\texttt{crd}(\overline{a}) \equiv^{\mathrm{qf}} \overline{a}'^\frown\texttt{crd}(\overline{a}')$.

The quartet $\mathbf{A} = (A, E, L, \texttt{crd})$ is called *partial self-coordinatized system* (pSCS) if A is flush, (L, E) is a $\sqsubseteq$-reflexive pair, crd satisfies the conditions just described and one additional condition:

5. Suppose $\overline{a}, \overline{a}_1, \overline{b}, \overline{b}_1 \in A$ such that $\overline{a}\overline{b} \equiv^{\mathrm{qf}} \overline{a}_1\overline{b}_1$, and suppose $(\overline{b}, \overline{e}) = \mathtt{crd}(\overline{a})$. Then, for all $\overline{e}_1 \in ||\mathfrak{M}||^{\leq m}$, if $\overline{a}\overline{b}\overline{e} \equiv^{\mathrm{qf}} \overline{a}_1\overline{b}_1\overline{e}_1$, then $\overline{e}_1 \in E$.

If $\mathbf{A} = (A, E, L, \mathtt{crd})$ is a pSCS, we define

$$\delta(\mathbf{A}) = \begin{cases} \max\left\{U^{\flat}(\overline{a}) : \overline{a} \in A \setminus dom(\mathtt{crd})\right\} & \text{if } A \setminus dom(\mathtt{crd}) \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

and we say that $\mathbf{A}$ is a self-coordinatized system (SCS) just in case $\delta(\mathbf{A}) = 0$. (Note that if $\mathbf{A}$ is not an SCS, then $\delta(\mathbf{A}) > 1$.)

We now consider the problem of enlarging a pSCS. Let $\mathbf{A} = (A, E, L, \mathtt{crd})$ be a pSCS, and suppose $\overline{a} \in A \setminus dom(\mathtt{crd})$ such that $U^{\flat}(\overline{a}) > 1$. Let $(\overline{b}, \overline{e})$ be a coordinate of $\overline{a}$ over $\emptyset$. Define $0 \leq \lambda_{\mathbf{A}}(\overline{b}\overline{e}) \leq \lceil r/t \rceil$ if there are $\overline{c}_1, ..., \overline{c}_n \in L$ such that

$$|rng(\overline{b}\overline{e}) \setminus (rng(\overline{c}_1) \cup \cdots \cup rng(\overline{c}_n))| \leq r$$

We say that $(\overline{b}, \overline{e})$ is an $\mathbf{A}$-*optimal coordinate* of $\overline{a}$ over $\emptyset$ if it minimizes $\lambda_{\mathbf{A}}(\overline{e})$ subject to the constraints $rng(\overline{b}) \cap acl(\overline{e}) = \emptyset$ and $rng(\overline{a}) \cap rng(\overline{b}) = \emptyset$. An *enlargement of* $\mathbf{A}$ is a pSCS $\mathbf{A}' = (A', E', L', \mathtt{crd}')$ such that $A \subseteq A'$, $E \subseteq E'$, $L \subseteq L'$ and $\mathtt{crd} \subseteq \mathtt{crd}'$.

We define the notion of a *locally optimal enlargement* of a given pSCS $\mathbf{A} = (A, E, L, \mathtt{crd})$. Firstly, if $\mathbf{A}$ happens to already be an SCS, then $\mathbf{A}$ is itself the unique locally optimal enlargement of $\mathbf{A}$. Now, assume that $\delta(\mathbf{A}) > 1$, and define a locally optimal enlargement $\mathbf{A}^+ = (A^+, E^+, L^+, \mathtt{crd}^+)$ as follows:

1. Choose $\overline{a} \in A \setminus dom(\mathtt{crd})$ to maximize $U^{\flat}(\overline{a})$, and choose an $\mathbf{A}$-optimal coordinate $(\overline{b}, \overline{e})$ of $\overline{a}$ over $\emptyset$.

2. Let $\lambda_{\mathbf{A}}(\overline{b}\overline{e}) = \lceil r/t \rceil$. If $r > 0$, choose $\overline{c}_1, ..., \overline{c}_n \in L$ such that

$$rng(\overline{b}\overline{e}) \setminus (rng(\overline{c}_1) \cup \cdots \cup rng(\overline{c}_n)) = \{x_1, ..., x_r\}$$

   Choose a minimal cover $X_1, ..., X_d$ of $[r]$ by t-multisets (so $d = \lceil r/t \rceil$) , and let $\overline{x}_j = (x_i : i \in X_j)$ for each $j \in [d]$.

3. Let

$$L^+ = L \cup \bigcup_{j=1}^{d} \left\{\overline{d} \in ||\mathfrak{M}||^{<\omega} : \overline{d} \equiv^{\mathrm{qf}} \overline{x}_j\right\}$$

$$A^+ = A \cup L^+ \cup \left\{\overline{b}' \in ||\mathfrak{M}||^{<\omega} : \overline{b}' \equiv^{\mathrm{qf}} \overline{b}\right\}$$

   and

$$E^+ = E \cup \bigcup_{j=1}^{d} \left\{\overline{d} \in ||\mathfrak{M}||^{<\omega} : \exists \overline{a}_1\overline{b}_1 \equiv^{\mathrm{qf}} \overline{a}\overline{b}.\ \overline{a}\overline{b}\overline{e} \equiv^{\mathrm{qf}} \overline{a}_1\overline{b}_1\overline{d}\right\}$$

4. Define $\mathtt{crd}^+ \supset \mathtt{crd}$ as necessary such that $(\overline{a}, (\overline{b}, \overline{e})) \in \mathtt{crd}^+$.

Under these definitions, the proof of the following proposition is more or less trivial.

Proposition 7.4. *Let* $\mathbf{A}$ *be a pSCS. Let* $\mathbf{B}_0 = \mathbf{A}$*, and for* $n < \omega$*, let* $\mathbf{B}_{n+1}$ *be a locally optimal enlargement of* $\mathbf{B}_n$*. Then, for some* $n_{\mathbf{A}} < \omega$*,* $\mathbf{B}_{n_{\mathbf{A}}}$ *is an SCS – i.e.* $\delta(\mathbf{B}_{n_{\mathbf{A}}}) = 0$*.*

Note that up to imposing a linear order on the (finite) set of quantifier-free t-types over $\emptyset$, the mapping $\mathbf{A} \mapsto \mathbf{A}^{opt} = \mathbf{B}_{n_{\mathbf{A}}}$ is well-defined. Consider an SCS $\mathbf{A} = (A, E, L, \mathtt{crd})$; let $\overline{a} \in A$ such that $U^{\flat}(\overline{a}) > 1$, and let $(\overline{b}, \overline{e}) = \mathtt{crd}(\overline{a})$. We, then, say that $\overline{b}$ is the $\mathbf{A}$-*successor of* $\overline{a}$ and that $\overline{e}$ is the $\mathbf{A}$-*support of the arrow* $\overline{a} \to \overline{b}$. (We also set $\overline{a} \to \star$ when $U^{\flat}(\overline{a}) \leq 1$.) Let $S_{\mathbf{A}} \subseteq (A \cup \{\star\}) \times (A \cup \{\star\})$ be the reflexive-transitive closure of the $\mathbf{A}$-successor relation. The $\mathbf{A}$-successor relation induces the structure of a rooted tree with vertex set $A \cup \{\star\}$, directed *towards* the root $\star$. If $\overline{a} \in A$ is not the $\mathbf{A}$-successor of an $\overline{a}_0 \in A$, then $\overline{a}$ is called an $\mathbf{A}$-leaf, which justifies the notation L in the definition of a pSCS.

We say that the SCS $\mathbf{A}$ is *insufficient* if there is a tuple $\overline{c} \in ||\mathfrak{M}||^t$ such that for every finite set X of $\mathbf{A}$-leaves, $\overline{c}$ is *not algebraic* over $\bigcup_{\overline{a} \in X} \left\{ rng(\overline{b}) : (\overline{a}, \overline{b}) \in S_{\mathbf{A}} \right\}$. In this case, we say that $\overline{c}$ is a witness to the insufficiency of $\mathbf{A}$. Clearly, if $\overline{c}$ is a witness to the insufficiency of $\mathbf{A}$ and $\overline{c}' \equiv^{\mathrm{qf}} \overline{c}$, then $\overline{c}'$ is also a witness, so without too much ambiguity, we may say that $qtp(\overline{c})$ is a witness as well. Of course, $\mathbf{A}$ is sufficient just in case it is not insufficient.

Now, suppose $\overline{c}$ is a witness to the insufficiency of an SCS $\mathbf{A}$. Define $\mathbf{A}{+}\overline{c} = (A', E, L', \mathtt{crd})$ so that $L' = L \cup \left\{ \overline{c}' \in ||\mathfrak{M}||^{<\omega} : \overline{c}' \equiv^{\mathrm{qf}} \overline{c} \right\}$ and $A' = A \cup L'$. Selecting $\overline{c}$ according to a fixed linear order on the set of quantifier-free t-types over $\emptyset$, the mapping $\mathbf{A} \mapsto \Upsilon(\mathbf{A})$ is well-defined, where

$$\Upsilon(\mathbf{A}) = \begin{cases} \mathbf{A} & \text{if } \mathbf{A} \text{ is sufficient} \\ (\mathbf{A}{+}\overline{c})^{opt} & \text{if } qtp(\overline{c}) \text{ is a ``minimal'' witness to insuff. of } \mathbf{A} \end{cases}$$

Under these definitions, the proof of the following proposition is, once again, essentially trivial.

Proposition 7.5. *Let* $\mathbf{A}$ *be a pSCS. Let* $\mathbf{B}_0 = (\mathbf{A})^{opt}$*, and for* $n < \omega$*, let* $\mathbf{B}_{n+1} = \Upsilon(\mathbf{B}_n)$*. Then, for some* $h_{\mathbf{A}} < \omega$*,* $\mathbf{B}_{h_{\mathbf{A}}}$ *is a sufficient SCS.*

If $\mathbf{A} = (A, E, L, \mathtt{crd})$ is SCS, we define $d_{\mathbf{A}} < \omega$ – *the stretch of* $\mathbf{A}$ – to be the smallest number $d < \omega$ such that for all $\overline{a} \in A$ such that $U^{\flat}(\overline{a}) > 1$, if $(\overline{b}, \overline{e}) = \mathtt{crd}(\overline{a})$, then there are $\overline{c}_1, ..., \overline{c}_d \in L$ such that $rng(\overline{b}\overline{e}) = rng(\overline{c}_1) \cup \cdots \cup rng(\overline{c}_d)$. Up to sacrificing optimality (after the fact), we may assume that for every $\overline{a} \in A$, $U^{\flat}(\overline{a}) > 1$, we have $\mathtt{crd}(\overline{a}) = (\overline{b}, \overline{e}_1{}^\frown \cdots {}^\frown \overline{e}_{d_{\mathbf{A}}})$ where $\overline{e}_i \in L$ for each $i = 1, ..., d_{\mathbf{A}}$. Abusing terminology somewhat, we then say that $\mathbf{A}$ has *constant stretch.*

7.3 A sufficient SCS as a structure

Throughout this section, let $\mathbf{A} = (A, E, L, \mathtt{crd})$ be a sufficient SCS with constant stretch $d_{\mathbf{A}}$. For some of the remaining analysis, it will be convenient to render $\mathbf{A}$ itself as a relational structure. The signatures of interest to us are:

$$\tau_{\texttt{crd}} = \left\{S_0^{(1)}, S_1^{(1)}, E^{(d_{\mathbf{A}})}, L^{(1)}, R_{\texttt{crd}}^{(2+d_{\mathbf{A}})}\right\} \cup \left\{R_{qtp(\overline{a})}^{(1+n)} : \overline{a} = (a_1, ..., a_n) \in A\right\}$$
$$\tau_{\texttt{crd}}^{\circ} = \left\{S_1^{(1)}, E^{(d_{\mathbf{A}})}, L^{(1)}, R_{\texttt{crd}}^{(2+d_{\mathbf{A}})}\right\} \cup \left\{R_{\circ,qtp(\overline{a})}^{(1)} : \overline{a} = (a_1, ..., a_n) \in A\right\}$$

We make $\mathbf{A}$ into a $\tau_{\texttt{crd}}$-structure with the following interpretations:

- $S_0^{\mathbf{A}} = ||\mathfrak{M}||$
- $S_1^{\mathbf{A}} = \{\eta(\overline{a}) : \overline{a} \in A\} \cup \{\star\}$ where $\eta : ||\mathfrak{M}||^{<\omega} \to \omega$ is an injective map. (Without loss of generality, we assume that $||\mathfrak{M}||^{<\omega} \cap \omega = \emptyset$.)
- $E^{\mathbf{A}} = \{(\eta(\overline{e}_1), ..., \eta(\overline{e}_d)) : \overline{e}_1{}^\frown \cdots {}^\frown \overline{e}_{d_{\mathbf{A}}} \in E\}$
- $L^{\mathbf{A}} = \{\eta(\overline{a}) : \overline{a} \in L\}$
- $R_{\texttt{crd}}^{\mathbf{A}}$ is the set of tuples $(\eta(\overline{a}), \eta(\overline{b}), \eta(\overline{e}'_1), ..., \eta(\overline{e}'_{d_{\mathbf{A}}}))$ such that
  1. $\overline{e}'_1, ..., \overline{e}'_{d_{\mathbf{A}}} \in L$
  2. $(\overline{b}, \overline{e}_1{}^\frown \cdots {}^\frown \overline{e}_{d_{\mathbf{A}}}) = \texttt{crd}(\overline{a})$ for some $\overline{e}_1, ..., \overline{e}_{d_{\mathbf{A}}} \in L$ such that

$$(\overline{a}, \overline{b}, \overline{e}'_1{}^\frown \cdots {}^\frown \overline{e}'_{d_{\mathbf{A}}}) \equiv^{\text{qf}} (\overline{a}, \overline{b}, \overline{e}_1{}^\frown \cdots {}^\frown \overline{e}_{d_{\mathbf{A}}})$$

- If $\sigma = qtp(\overline{a}_0)$, $\overline{a}_0 \in A$, then $R_{\sigma}^{\mathbf{A}} = \{(\eta(\overline{a}), a_1, ..., a_n) : \mathfrak{M} \models \sigma(\overline{a})\}$

Of course, $\mathbf{A}^{\circ}$ is defined in the same manner except that

$$R_{\circ,\sigma}^{\mathbf{A}^{\circ}} = \{\eta(\overline{a}) : \mathfrak{M} \models \sigma(\overline{a})\}$$

when $\sigma = qtp(\overline{a}_0)$, $\overline{a}_0 \in A$. The notation for $\mathbf{A}$ and $\mathbf{A}^{\circ}$ is slightly cumbersome as presented, so for economy in the sequel, we will often write R_{σ} when we mean $R_{\circ,\sigma}$. Naturally, there's no obstruction *per se* to making the same definitions for an insufficient SCS or even a pSCS. Finally, observe that for $X \subseteq ||\mathfrak{M}||$, $acl^{\mathbf{A}}(X) \subseteq acl^{\mathfrak{M}}(X)$, and if $Y \subseteq ||\mathbf{A}^{\circ}||$, then $acl^{\mathbf{A}^{\circ}}(Y) \subseteq acl^{\mathbf{A}}(Y)$

The first theorem, 7.6, of this section follows quite easily from the fact that $\mathfrak{M}$ is $\aleph_0$-categorical and (obviously) $\mathbf{A}$ is bi-interpretable (in fact, bi-definable) with $\mathfrak{M}$ without parameters. The corollary, 7.7, is a standard fact about countably categorical theories.

Theorem 7.6. *Let* $\mathbf{A} = (A, E, L, \texttt{crd})$ *be a sufficient SCS for* $\mathfrak{M}$*. Then as a* $\tau_{\texttt{crd}}$*-structure,* $\mathbf{A}$ *is* $\aleph_0$*-categorical, and* $\mathbf{A}^{\circ}$ *is* $\aleph_0$*-categorical as a* $\tau_{\texttt{crd}}^{\circ}$*-structure.*

Corollary 7.7. *Let* $\mathbf{A} = (A, E, L, \mathtt{crd})$ *be a sufficient SCS for* $\mathfrak{M}$*. Then, as structures,* $\mathbf{A}$ *and* $\mathbf{A}^\circ$ *are uniformly locally finite: There are functions* $g, g^\circ : \omega \to \omega$ *such that* $|acl^{\mathbf{A}}(X)| \leq g(|X|)$ *for any* $X \subset_{\mathit{fin}} ||\mathbf{A}||$ *and* $|acl^{\mathbf{A}^\circ}(Y)| \leq g^\circ(|Y|)$ *for any* $Y \subset_{\mathit{fin}} ||\mathbf{A}^\circ||$*.*

Substructures and localizations:

It will be convenient to designate a notion of substructure for SCS's and pSCS's; let $\mathbf{A} = (A, E, L, \mathtt{crd})$ and $\mathbf{A}_0 = (A_0, E_0, L_0, \mathtt{crd}_0)$ be pSCS's. . We write $\mathbf{A}_0 \leq \mathbf{A}$ just in case:

1. $A_0 \subseteq A$, $E_0 \subseteq E$ and $L_0 \subseteq L$
2. $R^{\mathbf{A}_0}_{\mathtt{crd}} = R^{\mathbf{A}}_{\mathtt{crd}} \cap (\eta[A_0] \times ((\eta[A_0] \times \eta[E_0]) \cup \{\star\}))$

We will also use a notion of localization for pSCS's. Specifically, suppose $\mathbf{A}$ is a pSCS and $X \subseteq ||\mathfrak{M}||$. We define $\mathbf{A}_X$ (as a structure) as follows:

- $S_0^{\mathbf{A}_X} = X$,
- $S_1^{\mathbf{A}_X} = S_1^{\mathbf{A}} \cap \{\eta(\overline{x}) : \overline{x} \in X^{<\omega}\} \cup \{\star\}$
  $E^{\mathbf{A}_X} = E^{\mathbf{A}} \cap (\{\eta(\overline{x}) : \overline{x} \in X^{<\omega}\})^d$
  $L^{\mathbf{A}_X} = L^{\mathbf{A}} \cap \{\eta(\overline{x}) : \overline{x} \in X^{<\omega}\}$
- $R_\sigma^{\mathbf{A}_X} = R_\sigma^{\mathbf{A}} \cap (\eta[X^{<\omega}] \times X^{<\omega})$
- $R_{\mathtt{crd}}^{\mathbf{A}_X} = R_{\mathtt{crd}}^{\mathcal{A}} \cap (\eta[X^{<\omega}])^{d+2}$

We say that $\mathbf{A}_X$ *is a localization of* $\mathbf{A}$ if for all $\eta(\overline{c}) \in S_1^{\mathbf{A}_X}$,

$$\mathbf{A} \models \exists y, z_1, ..., z_d(R_{\mathtt{crd}}(\eta(\overline{c}), y, \overline{z})) \;\Rightarrow\; \mathbf{A}_X \models \exists y, z_1, ..., z_d(R_{\mathtt{crd}}(\eta(\overline{c}), y, \overline{z}))$$

The project of this section, then, is to use the fact of uniform local finiteness to convert a sufficient SCS into a coherent solution of the model-building problem.

Surrogates for subsets of $||\mathfrak{M}||$:

We assume that $\mathbf{A} = (A, E, L, \mathtt{crd})$ is a sufficient SCS for $\mathfrak{M}$. Let $\tau = qtp(\overline{a}_0)$ for some $\overline{a}_0 \in A$.

$$V(\tau) = \left\{\mathcal{M} \in (Age(\mathfrak{M}) \cap K^G) : \exists \overline{a} \in \tau(\mathfrak{M}).\eta(\overline{a}) \in acl^{\mathbf{A}}(M)\right\} /\!\cong$$

We say that τ is *necessary* just in case $V(\tau)$ is a cofinite subset of $(Age(\mathfrak{M}) \cap K^G)/\!\cong$. If $C \subseteq S_1^{\mathbf{A}}$, then we set

$$\mathtt{proj}(C) = \bigcup \{rng(\overline{c}) : \eta(\overline{c}) \in C\}$$

Now, consider $X \subset ||\mathfrak{M}||$, and let $B \subseteq ||\mathbf{A}^\circ||$; we say that B is a *surrogate for* X *with respect to* $\mathbf{A}$ if

1. $X \cap fld(\eta^{-1}(S_1^{\mathbf{A}})) \subseteq \mathtt{proj}(B)$, and for $\overline{x} \in X^{<\omega}$, if $qtp(\overline{x})$ is necessary, then $\eta(\overline{x}) \in B$.
2. If τ is necessary, then there is an $\eta(\overline{a}) \in B$ such that $\mathfrak{M} \models \tau(\overline{a})$.
3. Suppose $\eta(\overline{a}_0), \eta(\overline{b}_0) \in B$ and $\overline{a}_1, \overline{b}_1 \in A$, $\overline{e}_1, ..., \overline{e}_d \in L$ such that $(\overline{b}_1, \overline{e}_1 \cdots \overline{e}_d) = \mathtt{crd}(\overline{a}_1)$ and $\overline{a}_1\overline{b}_1 \equiv^{\mathrm{qf}} \overline{a}_0\overline{b}_0$; let $\theta(\overline{x}, \overline{y}, \overline{z}) = qtp(\overline{a}_1, \overline{b}_1, \overline{e}_1 \cdots \overline{e}_d)$. Then there are $\eta(\overline{b}) \in B$ and $\eta(\overline{c}_1), ..., \eta(\overline{c}_d) \in B \cap L$ such that $\mathfrak{M} \models \theta(\overline{a}_0, \overline{b}, \overline{c}_1 \cdots \overline{c}_d)$.

The first lemma regarding surrogates is the following (whose proof we omit that we might avoid galling tedium).

Lemma 7.8 (Preparation lemma)**.** *There is a loop-free* `invent`*-program P_{surr}, depending on $\mathbf{A}^\circ$, such that (up to canonical embedding in $\mathbf{A}^\circ$) the follow holds: For every $X \subseteq ||\mathfrak{M}||$,* $\mathtt{resp}(P_{surr}, \mathfrak{M}{\upharpoonright}X)$ *is a surrogate for X with respect to $\mathbf{A}$. Furthermore, if X is finite, then* $\mathtt{resp}(P_{surr}, \mathfrak{M}{\upharpoonright}X)$ *is finite as well, and in any case,* $\mathtt{resp}(P_{surr}, \mathfrak{M}{\upharpoonright}X)$ *terminates in $|X|^{O(1)}$ many steps (when $\mathbf{A}$ is held constant).*

Proposition 7.9. *Let $X \subseteq ||\mathfrak{M}||$, and suppose $C \subseteq ||\mathbf{A}^\circ||$ is a surrogate for X with respect to $\mathbf{A}$. Let $D_1 = acl^{\mathbf{A}^\circ}(C)$, and let*

$$D_2 = \mathtt{resp}(\, P_{surr}, \mathfrak{M}{\upharpoonright}acl^{\mathfrak{M}}(\mathtt{proj}(D_1) \cup X)\,)$$

$$D = X \cup \mathtt{proj}(D_2)$$

Then there is a $\mathbf{B} \leq \mathbf{A}$, an SCS, such that $C \subseteq D_2 \subseteq ||\mathbf{B}^\circ||$ and $\mathbf{B}_D$ is a localization of $\mathbf{B}$.

In particular, if X is a finite set, then we may assume that C, D_1, D and $\mathbf{B}_D$ are each finite as well.

Sketch of the proof. Let $Q = \{qtp(\overline{a}) : \eta(\overline{a}) \in D_2\}$. Define $\mathbf{B} = \mathbf{B}[D_2]$ as follows:

- $S_0^{\mathbf{B}} = ||\mathfrak{M}||$
- $S_1^{\mathbf{B}} = \{\eta(\overline{a}) \in S_1^{\mathbf{A}} : qtp(\overline{a}) \in Q\} \cup \{\star\}$,
  $E^{\mathbf{B}} = E^{\mathbf{A}} \cap (S_1^{\mathbf{B}})^d$, $L^{\mathbf{B}} = L^{\mathbf{A}} \cap S_1^{\mathbf{B}}$.
- $R_\tau^{\mathbf{B}} = R_\tau^{\mathbf{A}}$ for all $\tau \in Q$ (and $R_\tau^{\mathbf{B}}$ does not exist if $\tau \notin Q$).
- $R_{\mathtt{crd}}^{\mathbf{B}} = R_{\mathtt{crd}}^{\mathbf{A}} \cap (S_1^{\mathbf{B}} \times S_1^{\mathbf{B}} \times E^{\mathbf{B}})$.

It's quite easy to see that $\mathbf{B} \leq \mathbf{A}$ and $C \subseteq D_2 \subseteq ||\mathbf{B}^\circ||$. It's also easy enough to see that D_2 is a surrogate of D with respect to $\mathbf{B}$ (even with respect to $\mathbf{A}$), and from this observation, it follows that $\mathbf{B}_D$ is a localization of $\mathbf{B}$. □

We observe that, up to the resolution of algebraic types, the transformation $X \rightsquigarrow \mathbf{B}_D$ in proposition 7.9 is computable in relation polynomial-time.

Proposition 7.10. *Assume $t > k$. Let $X \subseteq ||\mathfrak{M}||$ such that $X = acl^{\mathfrak{M}}(X)$ and $\mathbf{A}_X$ is a surrogate for X with respect to $\mathbf{A}$. As in proposition 7.9, let*

$$D_1(X) = acl^{\mathbf{A}^\circ}(\mathbf{A}_X)$$
$$D_2(X) = \mathtt{resp}(\,P_{surr}, \mathfrak{M}{\restriction}acl^{\mathfrak{M}}(\mathtt{proj}(D_1(X)) \cup X)\,)$$
$$D(X) = X \cup \mathtt{proj}(D_2(X))$$

Suppose:

1. $\mathbf{B} = \mathbf{B}[D_2(X)]$ *is a sufficient SCS for* $\mathfrak{M}$
2. $D(X) = X$
3. $\mathbf{B}_X$ *is a localization of* $\mathbf{B}$

Then $(\mathfrak{M}{\restriction}X)_{\rho^G}$ is a model of T^G.

Proof. We write X^G for the set $||(\mathfrak{M}{\restriction}X)_{\rho^G}||$ as well as the induced substructure with this universe. Since $\mathfrak{M}_{\rho^G}$ is model of T^G, we know that X^G satisfies the $\forall$-axioms G1-G4 of T^G (see chapter [not this one]). That X^G satisfies axiom G5 $= \bigwedge_\alpha \exists \overline{x}\, R_\alpha(\overline{x})$ follows immediately from the fact that $\mathbf{B}[D_2(X))]$ is sufficient for $\mathfrak{M}$. Thus, we need only verify that

$$X^G \vDash \bigwedge_{\alpha} \bigwedge_{\alpha \in Acc(\alpha,-)} \forall x_1...x_k(R_\alpha(\overline{x}) \to \exists y\, R_\beta(y, x_2, ..., x_k))$$

Let $(a_1, ..., a_k) \in R_\alpha^{X^G}$ and $\beta \in Acc(\alpha, -)$. Let τ_1,..., τ_n enumerate all quantifier-free t-types in the language of $\mathfrak{M}$ such that

$$\tau_i(\overline{x}, y, \overline{z}) \vDash \alpha(x_1, ..., x_k) \wedge \beta(y, x_2, ..., x_k)$$

By definition of sufficiency, we may assume that there are an $i \in [n]$ and a number $s < \omega$, for every $(\overline{a}', b', \overline{c}) \in \tau(\mathfrak{M})$, there are $\overline{d}_1, ..., \overline{d}_s \in L^{\mathbf{B}}$ such that $(\overline{a}', b, \overline{c})$ is algebraic over $Y(\overline{d}_1, ..., \overline{d}_s) = \bigcup_{i=1}^s \{rng(\overline{e}) : S^{\mathbf{B}}(\overline{d}_i, \overline{e})\}$. It follows that there is a formula $\varphi(y, \overline{a}')$ (not necessarily quantifier-free) which is uniformly algebraic in the type of $\overline{a}'$ and such that $\varphi(b', \overline{a}') \vDash R_\beta(b', a_2, ..., a_k')$. As X is algebraically-closed, there is a $b \in X$ such that $\varphi(b, \overline{a})$, so $R_\beta(b, a_1, ..., a_k)$, as required. $\square$

In light of the two preceding propositions, the next theorem is easy to verify:

Theorem 7.11. *By Propositions 7.9 and 7.10, the following procedure*

$$\mathtt{ComputeSurrogateSCSandExpand}(-; P_{surr}, \mathbf{A})$$

amounts to a coherent solution of the model-building problem for T.

Procedure: `ComputeSurrogateSCSandExpand`$(-; P_{surr}, \mathbf{A})$

1. Given $X \subset_{\mathbf{fin}} ||\mathfrak{M}||$
2. Set $V(X) = \bigcup_{i<\omega} U^i(X)$, where:
   (a) $U^0(X) = acl^{\mathfrak{M}}(X)$
   (b) For $i < \omega$, define
      i. $C^i(X) = \mathtt{resp}(P_{surr}, \mathfrak{M}{\restriction}U^i(X))$
      $D^i_1(X) = acl^{\mathbf{A}^\circ}(C^i(X))$
      $D^i_2(X) = \mathtt{resp}(\, P_{surr}, \mathfrak{M}{\restriction}acl^{\mathfrak{M}}(\mathtt{proj}(D^i_1(X)) \cup U^i(X))\,)$
      ii. $U^{i+1}(X) = U^i(X) \cup \mathtt{proj}(D^i_2(X))$
3. If $\mathtt{resp}(P_{surr}, \mathfrak{M}{\restriction}V(X))$ does not meet every type realized in $\mathbf{A}$, then add the optimal one, obtaining an extension $X' \supset X$, and return to 2 with X' in place of X.

Chapter 8

Unfolding digraphs and separation independence

Throughout this chapter, consider fully normalized `invent`-programs P over a signature ρ. That is, P is of the form $P_{pre}; P_{loop}; P_{post}$ – and more precisely of the form:

$$P_{pre}; (\ \textbf{while}\, \varphi_{loop}\, \textbf{do}\, (\textbf{if}\, \psi_1\, \textbf{then}\, R_1\, \textbf{else}\, \mathrm{id}); ...; (\textbf{if}\, \psi_m\, \textbf{then}\, R_m\, \textbf{else}\, \mathrm{id})\ ); P_{post}$$

satisfying the following conditions:

1. P_{pre} and P_{post} are sequences of basic expressions.
2. $\emptyset \vdash \bigvee_i \psi_i$ and if $i \neq j$, then $\psi_i \wedge \psi_j \vdash \textbf{false}$
3. For each $i = 1, ..., m$, R_i is a sequence of flat basic expressions.

We also assume that P is essentially inflationary.

8.1 Construction of the P-unfolding digraph

8.1.1 Construction of the naive R_i-unfolding digraphs

For this subsection (basically so that we can use the subscript i without worry), we fix a sequence of basic expressions $Q = \varepsilon_1; ...; \varepsilon_n$, but we keep the distinguished set of public variables $pub(P)$ in mind. Given a Γ-structure $\mathcal{A}$, we devise a digraph $G_Q^{\mathcal{A}}$ to represent the evaluation of $\texttt{resp}(Q, \mathcal{A})$.

As usual, we set $\mathcal{A}_0 = \mathcal{A}$ and $\mathcal{A}_{i+1} = \texttt{resp}(\varepsilon_{i+1}, \mathcal{A}_i)$ for $i = 1, ..., n$. Set $\texttt{var}(0) = \texttt{var}(n+1) = pub(P)$ and $\texttt{var}(i+1) = \texttt{var}(i) \cup \{head(\varepsilon_{i+1})\}$ if $i < n$. For $0 \leq i \leq n+1$, set

$$V_i(G_Q^{\mathcal{A}}) = \bigcup_{X \in \texttt{var}(i)} \begin{cases} \left\{(i, \overline{a}, X, +) : \overline{a} \in X^{\mathcal{A}_i}\right\} \\ \left\{(i, \overline{a}, X, -) : \overline{a} \in A_i^{ari(X)} \setminus X^{\mathcal{A}_i}\right\} \end{cases}$$

and set $V(G_Q^{\mathcal{A}}) = \bigcup_i V_i(G_Q^{\mathcal{A}})$. For $i < n$, set $E^0_{i,i+1}(G_Q^{\mathcal{A}}) =$

$$= \bigcup_{X \in \texttt{var}(i), X \neq head(\varepsilon_{i+1})} \begin{cases} \{((i,\overline{a},X,+),(i+1,\overline{a},X,+)) : \overline{a} \in X^{\mathcal{A}_i}\} \\ \{((i,\overline{a},X,-),(i+1,\overline{a},X,-)) : \overline{a} \in A^k_{i+1} \setminus X^{\mathcal{A}_i}\} \end{cases}$$

These pairs are called *trivial edges*, for they simply carry existing information forward when no change to the variable has occurred. By contrast, $E^1_{i,i+1}(G_Q^{\mathcal{A}})$ depends on ε_{i+1}:

- If $\varepsilon_{i+1} = (Y \leftarrow \{\overline{x} | \tau(\overline{x})\})$, then $E^1_{i,i+1}(G_Q^{\mathcal{A}}) = \emptyset$
- If $\varepsilon_{i+1} = (Y \leftarrow \{\overline{x} | X(\overline{x})\})$, then $E^1_{i,i+1}(G_Q^{\mathcal{A}}) =$

$$\{((i,\overline{a},X,+),(i+1,\overline{a},Y,+)) : \overline{a} \in X^{\mathcal{A}_i}\}$$

- If $\varepsilon_{i+1} = (Y \leftarrow \{\overline{x} | \neg X(\overline{x})\})$, then $E^1_{i,i+1}(G_Q^{\mathcal{A}}) =$

$$\{((i,\overline{a},X,-),(i+1,\overline{a},Y,+)) : \overline{a} \in A^k_i \setminus X^{\mathcal{A}_i}\}$$

- If $\varepsilon_{i+1} = (Y \leftarrow \{\overline{x} | \exists \overline{y}\, X(\overline{x},\overline{y})\})$, then $E^1_{i,i+1}(G_Q^{\mathcal{A}}) =$

$$\{\big((i,\overline{a}\overline{b},X,+),(i+1,\overline{a},Y,+)\big) : \overline{a}\overline{b} \in X^{\mathcal{A}_i}\}$$

- If $\varepsilon_{i+1} = (Y \leftarrow \{\overline{x} | X_1(\overline{x}) \wedge X_2(\overline{x}_2)\})$, then $E^1_{i,i+1}(G_Q^{\mathcal{A}}) =$

$$\{((i,\overline{a}\restriction\overline{x}_j,X_j,+),(i+1,\overline{a},Y,+)) : \overline{a} \in A^k_i,\ \overline{a}\restriction\overline{x}_j \in X^{\mathcal{A}_i},\ j=1,2\}$$

- If $\varepsilon_{i+1} = (Y \leftarrow \{\overline{x} | X_1(\overline{x}) \vee X_2(\overline{x}_2)\})$, then $E^1_{i,i+1}(G_Q^{\mathcal{A}}) =$

$$= \bigcup \begin{cases} \{((i,\overline{a}\restriction\overline{x}_1,X_1,+),(i+1,\overline{a},Y,+)) : \overline{a} \in A^k_i,\ \overline{a}\restriction\overline{x}_1 \in X^{\mathcal{A}_1}, \overline{a}\restriction\overline{x}_2 \notin X^{\mathcal{A}_2}\} \\ \{((i,\overline{a}\restriction\overline{x}_2,X_2,+),(i+1,\overline{a},Y,+)) : \overline{a} \in A^k_i,\ \overline{a}\restriction\overline{x}_1 \notin X^{\mathcal{A}_1}, \overline{a}\restriction\overline{x}_2 \in X^{\mathcal{A}_2}\} \\ \{((i,\overline{a}\restriction\overline{x}_j,X_j,+),(i+1,\overline{a},Y,+)) : \overline{a} \in A^k_i,\ \overline{a}\restriction\overline{x}_j \in X^{\mathcal{A}_j}, j=1,2\} \end{cases}$$

- If $\varepsilon_{i+1} = (Y \leftarrow \texttt{invent}_k\, \{x_1 ... x_s : X(\overline{x})\})$, $s < k$, then $E^1_{i,i+1}(G_Q^{\mathcal{A}}) =$

$$\{((i,\overline{a},X,+),(i+1,\texttt{invent}_k(\overline{a}),Y,+)) : \overline{a} \in X^{\mathcal{A}_i}\}$$

  where for $\overline{a} \in A^s_i$, we define

$$\texttt{invent}_k(\overline{a}) = (a_1, ..., a_s, 1\widehat{\ }\overline{a}, ..., (k-s)\widehat{\ }\overline{a})$$

Set $E_{i,i+1}(G_Q^{\mathcal{A}}) = E^0_{i,i+1}(G_Q^{\mathcal{A}}) \cup E^1_{i,i+1}(G_Q^{\mathcal{A}})$, and $E(G_Q^{\mathcal{A}}) = \bigcup_{i \leq n} E_{i,i+1}(G_Q^{\mathcal{A}})$.

8.1.2 Extension to the P_{loop}-digraph and to the P-digraph

For convenience in typesetting, we now give the inner part of P_{loop} the name Q:

$$(\mathbf{if}\,\psi_1\,\mathbf{then}\,R_1\,\mathbf{else}\,\mathrm{id});...;(\mathbf{if}\,\psi_m\,\mathbf{then}\,R_m\,\mathbf{else}\,\mathrm{id})$$

If $\mathcal{A}$ is a ρ-structure such that $\texttt{resp}(Q,\mathcal{A})$ is defined, then by definition, $\texttt{resp}(Q,\mathcal{A}) = \mathcal{A}^{(n^*)}$ where

$$\begin{aligned}\mathcal{A}^{(0)} &= \mathcal{A}\\ \mathcal{A}^{(i+1)} &= \texttt{resp}(Q,\mathcal{A}^{(i)})\\ n^* &= \min\left\{i<\omega : \mathcal{A}^{(i)} \vDash \neg\varphi_{loop}\right\}\end{aligned}$$

Furthermore, there is a well-defined map

$$e = e_{\mathcal{A}} : \{0,1,...,n^*-1\} \longrightarrow \{1,...,m\}$$

such that $e(i) = t$ if and only if $\mathcal{A}_i \vDash \psi_t$. As previously, for each $i \in \{1,...,m\}$, l_i denotes the length of R_i – that is, the number of basic expressions comprising $R_i = \varepsilon_1^i;...;\varepsilon_{l_i}^i$. Thus, for a ρ-structure $\mathcal{A}$, the levels of $G_{R_i}^{\mathcal{A}}$ are numbered $0,1,...,l_i,l_i+1$.

We obtain the Q-digraph by connecting $G_{R_{e(0)}}^{\mathcal{A}^{(0)}}$ to $G_{R_{e(1)}}^{\mathcal{A}^{(1)}}$ to $G_{R_{e(2)}}^{\mathcal{A}^{(2)}}$ and so forth. More formally,

$$V(G_{P_{loop}}^{\mathcal{A}}) = \bigcup_{i\leq n^*}\left(\{i\}\times V(G_{R_{e(i)}}^{\mathcal{A}^{(i)}})\right)$$

$$E^{(i)}(G_Q^{\mathcal{A}}) = \left\{((i,u),(i,v)) : (u,v)\in E(G_{R_{e(i)}}^{\mathcal{A}^{(i)}})\right\}$$

$$E^{(i,i+1)}(G_Q^{\mathcal{A}}) = \cdots$$

$$\cdots = \bigcup_{X\in pub(P)}\left\{\big((i,(l_{e(i)},\overline{a},X,\pm)),(i+1,(0,\overline{a},X,\pm))\big) : (l_{e(i)},\overline{a},X,\pm)\in V_{l_{e(i)}}(G_{R_{e(i)}}^{\mathcal{A}^{(i)}})\right\}$$

$$E(G_Q^{\mathcal{A}}) = \bigcup_{i\leq n^*}E^{(i)}(G_Q^{\mathcal{A}}) \cup \bigcup_{i<n^*}E^{(i,i+1)}(G_Q^{\mathcal{A}})$$

The subsequent definition of $G_P^{\mathcal{A}}$ is much more concise. Setting

$$\mathcal{A}_{pre} = \texttt{resp}(P_{pre},\mathcal{A})\restriction pub(P_{loop})$$

we then define $G_P^{\mathcal{A}} = G_{P_{loop}}^{\mathcal{A}_{pre}}$

8.1.3 Pruning by algebraicity

The naive P_{loop}-unfolding digraphs, $G^{\mathcal{A}}_{P_{loop}}$, are really quite naive, as the name suggests; in a sense they retain *too much* information about the computation. To dispel this difficulty, we prune away much of the naive P_{loop}-unfolding digraph in a manner that leaves (what seems to be) only the model-theoretically interesting data. In particular, only public (hence, explicitly inflationary) relation variables are retained.

If K is a class of $pub(P)$-structures, we set

$$K_{pre} = \{\mathtt{resp}(P_{pre}, \mathcal{A}) : \mathcal{A} \in K\}$$

$$\mathtt{resp}(P, K)^* = \{\mathcal{A}_i : \mathcal{A} \in K, i = 0, 1, ..., n_{\mathcal{A}}\}$$

where for $\mathcal{A} \in K$, we define $n_{\mathcal{A}} < \omega$ so that

$$\begin{aligned} \mathcal{A}^{(0)} &= \mathcal{A}_{pre} \\ \mathcal{A}^{(i+1)} &= \mathtt{resp}(R, \mathcal{A}^{(i)}) \\ n_{\mathcal{A}} &= \min \left\{ i < \omega : \mathcal{A}^{(i)} \vDash \neg\varphi_{loop} \right\} \end{aligned}$$

Finally, we set

$$K^P = \{\mathcal{M} \restriction pub(P) : \mathcal{M} \in \mathtt{resp}(P, K)^*\}$$

For $\mathcal{M} \in K^P$ and $A \subseteq M$, we define (as usual),

$$K^P_A = K^P_{(A;\mathcal{M})} = \left\{ \mathcal{N} \in K^P : A \subseteq N,\ diag^{\mathcal{N}}(A) = diag^{\mathcal{M}}(A) \right\}$$

Suppose $m, n \in \{ari(X) : X \in pub(P)\}$, and let $p(x_1, ..., x_m, y_1, ..., y_n)$ be a quantifier- free type over $\emptyset$ in the language of $pub(P)$, which we assume is complete. We say that $p(\overline{x}, \overline{y})$ is an *algebraic step* with respect to K^P just in case it is

1. $p(\overline{x}, \overline{y})$ itself is non-algebraic over K^P: for every $r < \omega$, there is a model $\mathcal{M} \in K^P$ such that
$$\left| \left\{ \overline{a}\overline{b} \in M^{m+n} : \mathcal{M} \vDash p(\overline{a}, \overline{b}) \right\} \right| \geq r$$
2. There is a number $0 < r < \omega$ such that for all $\mathcal{M} \in K^P$ and $\overline{a} \in M^m$,
$$\left| \left\{ \overline{b} \in M^n : \mathcal{M} \vDash p(\overline{a}, \overline{b}) \right\} \right| \leq r$$

We are now prepared to construct the true P-unfolding digraph $H^{\mathcal{A}}_P$ for a structure $\mathcal{A} \in K$. Again, we have $\mathtt{resp}(P_{loop}, \mathcal{A}_{pre}) = \mathcal{A}^{(n_{\mathcal{A}})}$. For the vertex set, have simply

$$V(H^{\mathcal{A}}_P) = \bigcup_{0 \leq i \leq n_{\mathcal{A}}} \left\{ (i, (0, \overline{a}, X, \pm)) \in V(G^{\mathcal{A}}_P) : X \in pub(P) \right\}$$

As in the naive construction, we define edges only between adjacent levels – that is, we define the edge sets $E_{i,i+1}(H^{\mathcal{A}}_P)$. For each $i < n_{\mathcal{A}}$, $E_{i,i+1}(H^{\mathcal{A}}_P)$ contains two kinds of edges:

1. If $(i,(0,\overline{a},X,\xi))$ and $(i+1,(0,\overline{a},X,\xi))$, $\xi=\pm$ are in $V(H_P^{\mathcal{A}})$, then
$$\langle\, (i,(0,\overline{a},X,\xi)),(i+1,(0,\overline{a},X,\xi))\,\rangle$$
is in $E_{i,i+1}(H_P^{\mathcal{A}})$.

2. Suppose $(i,(0,\overline{a},X,\xi))$ and $(i+1,(0,\overline{b},Y,+))$ are in $V(H_P^{\mathcal{A}})$, where $(0,\overline{a},X,\xi)$ and $(0,\overline{b},Y,+)$ are different from each other and $\xi\in\{+,-\}$. Then
$$\Big\langle\, (i,(0,\overline{a},X,\xi)),(i+1,(0,\overline{b},Y,+))\,\Big\rangle$$
is in $E_{i,i+1}(H_P^{\mathcal{A}})$ just in case:
   - There is a directed path from $(i,(0,\overline{a},X,\xi))$ to $(i+1,(0,\overline{b},Y,+))$ in the naive unfolding graph $G_P^{\mathcal{A}}$;
   - The quantifier-free type $p(\overline{x},\overline{y})$ of $(\overline{a},\overline{b})$ in $\mathcal{A}^{(i+1)}$ is an algebraic step.
   - $(i,(0,\overline{b},Y,+))$ is *not in* $V(H_P^{\mathcal{A}})$

8.2 Definitions towards an independence relation

Digraphs and d-separation:

Let $G=(V,E)$ be a directed graph (digraph). For $u,v\in V$, we say that v is a descendant of u if there is a (directed) path from u to v in G; we denote by $Desc(u)$ the set containing u and all of u's descendants in G. A *trail* in G is a sequence of vertices $t=(v_1,...,v_n)$ such that for each $i=1,...,n-1$, either (v_i,v_{i+1}) or (v_{i+1},v_i) is an edge – that is, a trail is a path in the undirected graph associated with G. The internal vertices of t, $v_2,...,v_{n-1}$, can be classified into three types:

1. v_i is *head-to-head* if (v_{i-1},v_i) and (v_{i+1},v_i) are both edges.

2. v_i is *tail-to-tail* if (v_i,v_{i-1}) and (v_i,v_{i+1}) are both edges.

3. v_i is *head-to-tail* just in case either (v_{i-1},v_i) and (v_i,v_{i+1}) are edges, or (v_{i+1},v_i) and (v_i,v_{i-1}) are edges.

We can now define the key notion of *d-separation* in digraphs. Firstly, suppose A,B,C are pairwise disjoint subsets of V; we say that A and B are d-separated by C in G if for each trail $t=(v_1,...,v_n)$ with $v_1\in A$ and $v_n\in B$, at least one of the following conditions holds for some $1<i<n$:

1. v_i is tail-to-tail in t and $v_i\in C$

2. v_i is head-to-tail in t and $v_i\in C$

3. v_i is head-to-head in t and $Desc(v_i) \cap C = \emptyset$.

Extending this notion, if A, B, C are not necessarily pairwise disjoint, we again say that A and B are d-separated by C just in case (a) $A \cap B \subseteq C$ and (b) $A \setminus C$ and $B \setminus C$ are d-separated by C. We use the notation $[A \amalg B \,|\, C]_G$ to abbreviate the assertion, "A and B are d-separated by C in G". The following properties of the d-separation relation are not difficult to derive, so their proofs are omitted. In any case, the arguments can be found in [26] or [4], and the terminology (which we consider somewhat unfortunate) is taken from those sources:

Lemma 8.1. *Let $G = (V, E)$ be a directed acyclic graph. The relation of d-separation has the following properties:*

1. Symmetry: $[A \amalg B \,|\, C]_G \;\Rightarrow\; [B \amalg A \,|\, C]_G$
2. Decomposition: $[A \amalg B \,|\, C]_G \wedge B_0 \subseteq B \;\Rightarrow\; [A \amalg B_0 \,|\, C]_G$
3. Weak Union: $[A \amalg B \,|\, C]_G \wedge B_0 \subseteq B \;\Rightarrow\; [A \amalg B \,|\, C \cup B_0]_G$
4. Contraction:
   $[A \amalg B_1 \,|\, C]_G \wedge [A \amalg B_2 \,|\, C \cup B_1]_G \;\Rightarrow\; [A \amalg B_1 \cup B_2 \,|\, C]_G$

Algebraic closure, local separation and deviation:

We will assume that the set of (situated) structures K under consideration is closed under isomorphism and has the amalgamation property with respect to quantifier-free formulas. Since we don't have access to quantifiers, it's necessary to define the following closure operator (which we actually defined in the Introduction) by an explicit recursion:

$$\begin{aligned}
\kappa^{\mathcal{M}}(A) &= A \cup \{b \in M : qtp(b/A; \mathcal{M}) \text{ is } K\text{-algebraic}\} \\
\mathrm{cl}_0(A; \mathcal{M}) &= A \\
\mathrm{cl}_{n+1}(A; \mathcal{M}) &= \kappa^{\mathcal{M}}(acl_n(A; \mathcal{M})) \\
\mathrm{cl}(A; \mathcal{M}) &= \mathrm{cl}_{|M|}(A; \mathcal{M})
\end{aligned}$$

We write $\mathrm{cl}(A)$ in place of $\mathrm{cl}(A; \mathcal{M})$ when $\mathcal{M}$ is clear from context. If $\mathcal{M} \in K$ and $A \subseteq M$, then we identify A with the subset of vertices

$$V(H_P^{\mathcal{M}}) \cap \left\{(0, (0, \overline{a}, X, \pm)) : X \in pub(P),\, \overline{a} \in A^{ari(X)}\right\}$$

We also define a technical notion of *hereditary descendants*, $HDesc(A; \mathcal{M})$, which keeps closer to the original base set in question. First, define

$$HDesc^0(A; \mathcal{M}) = \left\{(i, (0, \overline{a}, X, \pm)) \in V(H_P^{\mathcal{M}}) : \overline{a} \in \mathrm{HL}[A]\right\}$$

Next, suppose $(i, v) \in HDesc^0(A; \mathcal{M})$.

- If $v = (0, \overline{a}, X, +)$, then $(i, v) \in HDesc(A; \mathcal{M})$
- If $v = (0, \overline{a}, X, -)$, then $(i, v) \in HDesc(A; \mathcal{M})$ only if there is a $(j, w) \in HDesc^0(A; \mathcal{M})$ such that $w \neq v$ (so $sign(w) = +$), $i < j$, and $(j, w) \in Desc_{H_P^{\mathcal{M}}}(A)$.

Again, we sometimes write $HDesc(A)$ in place of $HDesc(A; \mathcal{M})$ when the structure $\mathcal{M}$ is clear from context. Clearly, $HDesc(A) \subseteq HDesc(B)$ when $A \subseteq B \subseteq M$.

Finally, we define the *local separation* relation on subsets of K-models as follows: Let $\mathcal{M} \in K$ and $A, B, C \subseteq M$; then we write

$$A \underset{C}{\downarrow} B/\mathcal{M}$$

– and say that A *and* B *are locally separated by* C *in* $\mathcal{M}$ – just in case

$$[\, A \amalg B \,|\, HDesc(acl(D)) \,]_{H_P^{\mathcal{M}}}$$

whenever $\mathrm{cl}(C; \mathcal{M}) \subseteq D \subseteq \mathrm{cl}(BC; \mathcal{M})$.

Lemma 8.2 (Local monotonicities). *Let $\mathcal{M} \in K$, $A, B, C \subseteq M$ and $B_0 \subseteq B$. If $A \downarrow_C B/\mathcal{M}$, then*

1. *(Monotonicity) $A \downarrow_C B_0/\mathcal{M}$*
2. *(Base-monotonicity) $A \downarrow_{CB_0} B/\mathcal{M}$*

Proof. 1. Suppose $acl(C) \subseteq D \subseteq acl(B_0C)$. Since $acl(B_0C) \subseteq acl(BC)$ and $A \downarrow_C B/\mathcal{M}$, we have

$$[A \amalg B \,|\, HDesc(\mathrm{cl}(D))]_{H_P^{\mathcal{M}}}$$

and by Decomposition, we have

$$[A \amalg B_0 \,|\, HDesc(\mathrm{cl}(D))]_{H_P^{\mathcal{M}}}$$

As D was arbitrary, it follows that $A \downarrow_C B_0/\mathcal{M}$.

2. Suppose $acl(B_0C) \subseteq D \subseteq acl(B \cup (B_0C)) = acl(BC)$. Since $acl(C) \subseteq acl(B_0C)$ and $A \downarrow_C B/\mathcal{M}$, we have

$$[A \amalg B \,|\, HDesc(\mathrm{cl}(D))]_{H_P^{\mathcal{M}}}$$

Again, as D was arbitrary, we have shown that $A \downarrow_{B_0C} B/\mathcal{M}$. □

We will say that a partial (quantifier-free) type $\pi(x_1, ..., x_n)$ over a set $A \subseteq M$, where $\mathcal{M} \in K$, is *admissible* if $diag(A; \mathcal{M}) \subseteq \pi$; that is, π itself fixes the quantifier-free type of its own domain. Given an admissible type $\pi(\overline{x})$ over A, we write K_A^π for the set of all $\mathcal{N} \in K_A$

such that π is realized in $\mathcal{N}$. By the amalgamation assumption, K_A^π is infinite whenever π is K-consistent.

Suppose $\mathcal{M}_0 \in K$, $B_0 \subseteq B \subseteq M_0$, and let $\pi(x_1, ..., x_n)$ be an admissible type over B. Suppose $\mathcal{M} \in K_B$ and $B \subseteq D \subseteq M$. We define $\Delta_0^{\mathcal{M}}[\pi, B_0](D)$ to be the set of all structure $\mathcal{N} \in K_D^\pi$ such that for all $\overline{a} \in N^n$, if $\mathcal{N} \vDash \pi(\overline{a})$, then

$$\overline{a} \downarrow_B D/\mathcal{N}' \;\Rightarrow\; \overline{a} \not\downarrow_{B_0} D/\mathcal{N}'$$

whenever $\mathcal{N}' \in K_{D\overline{a}}$.

Note that if $\mathcal{N} \in \Delta_0^{\mathcal{M}}[\pi, B_0](D)$ and $\mathcal{N}' \cong_D \mathcal{N}$, then $\mathcal{N}' \in \Delta_0^{\mathcal{M}}[\pi, B_0](D)$. Hence,

$$\Delta^{\mathcal{M}}[\pi, B_0](D) = \Delta_0^{\mathcal{M}}[\pi, B_0](D)/\!\cong_D \subseteq K_D^\pi/\!\cong_D$$

is well-defined, and we call this set of isomorphism types the *deviation of* π *over* B_0 *with respect to* D.

The notions of freeness and independence:

It is standard practice of the first-order model theory of infinite structures (more or less) to identify all finite numbers with each other and understand all finite objects as “rank 0.” Following this practice, we will understand the statement, “ the deviation of π over B_0 with respect to D is finite,” as equivalent to, “π has null deviation over B_0 ,” or ”π does not deviate over B_0 with respect to D.”

It will be easiest – at least initially – to express our notion of independence as a set I_P of pairs (π, B_0) in which $\pi(x_1, ..., x_n)$ is an admissible type over some set $B \supseteq B_0$. Namely, we define I_P to be just the subset of such pairs (π, B_0) such that for all $\mathcal{M} \in K_B$, $B = dom(\pi)$, and $B \subseteq D \subseteq M$, $\Delta^{\mathcal{M}}[\pi, B_0](D)$ is finite. I_P, we say, is the *weak notion of freeness* induced by P. Furthermore, if $\pi(\overline{x})$ is an admissible type over B, $B_0 \subseteq B$, $\pi_0 = \pi \restriction B_0 \cup diag(B)$ is the *admissible restriction* of π to B, and $(\pi, B_0) \in I_P$, then we say that π *is a non-deviating extension of* π *to* B.[1] The *notion of independence* induced by P, then, is given by

$$\overline{a} \mathop{\downfall}^{\mathrm{d}}_C B \;\Leftrightarrow\; (qtp(\overline{a}/BC), C) \in I_P$$

when $\mathcal{M} \in K$, $B, C \subseteq M$ and $\overline{a} \in M^{<\omega}$. In the next few sections, we get down to proving that this notion of independence is indeed an independence relation as defined in the Introduction.

Author’s note on the definitions: Firstly, it seems quite likely that the sets $\Delta^{\mathcal{M}}[\pi, C](D)$ will always be either empty or equal to $K_D^\pi/\!\cong_D$ – and that is the intuition behind the definition – but I haven’t been able to prove this fact. Thankfully, the stronger is not necessary for the development. Secondly, some readers have found the definition of the expression $A \downarrow_C B/\mathcal{M}$ somewhat strangely formed insofar as the ranging parameter $\mathrm{cl}(C; \mathcal{M}) \subseteq D \subseteq \mathrm{cl}(BC; \mathcal{M})$

[1] With heartfelt apologies to francophones.

seems, at first look, unnecessarily complicated. However, it seems to me implausible that the local base-monotonicity property holds without this complication. Specifically, in the scenario, $\mathrm{cl}(C;\mathcal{M}) \subseteq D \subseteq \mathrm{cl}(BC;\mathcal{M})$, one could find that

$$[\, A \amalg B \mid HDesc(\mathrm{cl}(D;\mathcal{M}))\,]_{H_P^{\mathcal{M}}}$$

holds but

$$[\, A \amalg B \mid HDesc(\mathrm{cl}(BC;\mathcal{M}))\,]_{H_P^{\mathcal{M}}}$$

via a head-to-head collision with a descendent in $HDesc(\mathrm{cl}(BC;\mathcal{M}))$. In fact, this complexity, I believe, is the reason why we should not expect the independence relation to collapse to algebraic closure in general.

8.3 Basic properties of the notion of independence

In this section, we prove properties of our notion independence which are absolute in the sense that we need not make any new assumptions about the class K or the program P beyond those already made. That is, K has the amalgamation property, and P is a fully normalized essentially inflationary program which terminates on every structure in K. The reader will likely note that these basic properties are the same as those that hold of forking independence and þ-forking independence in an arbitrary theory – that is, in a theory that is not necessarily simple or rosy, respectively.

Lemma 8.3. *Let $\mathcal{M}_0 \in K$ and $B, C \subseteq M_0$, and let $\pi(\overline{x})$ be an admissible type over BC.*

1. *Suppose $\mathcal{M} \in K_{BC}$ and $BC \subseteq D_0 \subseteq D \subseteq M$. If $\mathcal{N} \in K_D^{\pi} \setminus \Delta_0^{\mathcal{M}}[\pi, C](D)$, then $\mathcal{N} \in K_{D_0}^{\pi} \setminus \Delta_0^{\mathcal{M}}[\pi, C](D_0)$.*
2. *Suppose $\mathcal{M} \in K_{BC}$ and $BC \subseteq D \subseteq M$, and suppose $\pi_0 \subseteq \pi$ is an admissible type over BC. If $\mathcal{N} \in K_D^{\pi} \setminus \Delta_0^{\mathcal{M}}[\pi, C](D)$, then $\mathcal{N} \in K_D^{\pi_0} \setminus \Delta_0^{\mathcal{M}}[\pi_0, C](D)$.*

Proof. 1. Let $\mathcal{N} \in K_D^{\pi} \setminus \Delta_0^{\mathcal{M}}[\pi, C](D)$. Suppose $\mathcal{N}' \in K_{D\overline{a}}$, $\overline{a} \downarrow_{BC} D/\mathcal{N}'$ and $\overline{a} \downarrow_C D/\mathcal{N}'$ for some $\overline{a} \in \pi(N)$. Then $\overline{a} \downarrow_{BC} D_0/\mathcal{N}'$ and $\overline{a} \downarrow_C D_0/\mathcal{N}'$ by Local monotonicity, and the claim follows.

2. Again, suppose $\overline{a} \downarrow_{BC} D/\mathcal{N}'$ and $\overline{a} \downarrow_C D/\mathcal{N}'$ for some $\overline{a} \in \pi(N)$. Clearly, $\overline{a} \in \pi_0(N)$, and the claim follows. □

Lemma 8.4 (I_P-Existence). *Suppose $\mathcal{M}_0 \in K$, $B \subseteq M_0$, and $\pi(\overline{x})$ is an admissible type over B. Then $(\pi, B) \in I_P$.*

Proof. Let $\mathcal{M} \in K_B$ and $B \subseteq D \subseteq M$, and suppose $\mathcal{N} \in K_D^{\pi}$. If $\mathcal{N} \models \pi(\overline{a})$ and $\overline{a} \downarrow_B D/\mathcal{N}$, then obviously the assertion $\overline{a} \not\downarrow_B D/\mathcal{N}$ is nonsense, so $\mathcal{N}$ is not in $\Delta_0^{\mathcal{M}}[\pi, B](D) = \emptyset$. $K_D^{\pi}/\cong_D$ is infinite by the amalgamation assumption, so $\Delta^{\mathcal{M}}[\pi, B](D)$ is a coinfinite subset, as required. □

Lemma 8.5 (I_P-Monotonicity). *Let $\mathcal{M}_0 \in K$ and $B, C \subseteq M_0$, and let $\pi(\overline{x})$ be an admissible type over BC such that $\pi{\restriction}B_0C$ is a complete quantifier-free type over B_0C in $\mathcal{M}_0$. If $(\pi, C) \in I_P$ and $B_0 \subseteq B$, then $(\pi{\restriction}B_0C, C) \in I_P$.*

Proof. Let $\mathcal{M} \in K_{B_0C}$ and $B_0C \subseteq D \subseteq M$. By the amalgamation property, we may choose a structure $\mathcal{N}_1 \in K^{\pi}_{BD}$ such that $\mathcal{N} \leq \mathcal{N}_1$. If $\mathcal{N}_1 \notin \Delta_0^{\mathcal{M}}[\pi, C](D)$, then we may choose $\overline{a} \in \pi(N_1)$ and $\mathcal{N}' \in K^{\pi}_{BD\overline{a}}$ such that $\overline{a} \downarrow_{BC} D/\mathcal{N}'$ and $\overline{a} \downarrow_{C} D/\mathcal{N}'$. This shows that that $\mathcal{N} \notin \Delta_0^{\mathcal{M}}[\pi{\restriction}B_0C, C](D)$ *provided that* $qtp(\overline{a}/B_0C; \mathcal{N}_1)$ is realized in $\mathcal{N}$. By hypothesis, $\pi{\restriction}B_0C = qtp(\overline{a}/B_0C; \mathcal{N}_1)$, and the lemma follows. □

Lemma 8.6 (I_P-Base-monotonicity). *Let $\mathcal{M}_0 \in K$ and $B, C \subseteq M_0$, and let $\pi(\overline{x})$ be an admissible type over BC. If $(\pi, C) \in I_P$ and $B_0 \subseteq B$, then $(\pi, B_0C) \in I_P$.*

Proof. Let $\mathcal{M} \in K_{BC}$ and $BC \subseteq D \subseteq M$. Suppose $\mathcal{N} \in K^{\pi}_{D} \setminus \Delta_0^{\mathcal{M}}[\pi, C](D)$. We may, then, choose $\overline{a} \in N^{\overline{x}}$ be such that $\mathcal{N} \vDash \pi(\overline{a})$, $\overline{a} \downarrow_{BC} D/\mathcal{N}$ and $\overline{a} \downarrow_{C} D/\mathcal{N}$ all hold. By local base-monotonicity, the last of these facts implies $\overline{a} \downarrow_{B_0C} D/\mathcal{N}$, so $\overline{a}$ witnesses the fact that $\mathcal{N} \in K^{\pi}_{D} \setminus \Delta_0^{\mathcal{M}}[\pi, B_0C](D)$. Thus,

$$K^{\pi}_{D} \setminus \Delta_0^{\mathcal{M}}[\pi, C](D) \;\subseteq\; K^{\pi}_{D} \setminus \Delta_0^{\mathcal{M}}[\pi, B_0C](D)$$

and it follows that

$$K^{\pi}_{D}/{\cong_D} \setminus \Delta^{\mathcal{M}}[\pi, C](D) \;\subseteq\; K^{\pi}_{D}/{\cong_D} \setminus \Delta^{\mathcal{M}}[\pi, B_0C](D)$$

and therefore,

$$\Delta^{\mathcal{M}}[\pi, B_0C](D) \subseteq \Delta^{\mathcal{M}}[\pi, C](D)$$

Since the term on the right-hand side is a finite set, we know that $\Delta^{\mathcal{M}}[\pi, B_0C](D)$ is finite, and as $\mathcal{M}, D$ were arbitrary, the proof of the lemma is complete. □

Proposition 8.7 (I_P-Extension). *Let $\mathcal{M}_0 \in K$ and $B, C \subseteq M_0$, and let $\pi(\overline{x})$ be an admissible type over BC. Suppose $(\pi, C) \in I_P$. Then there is a complete extension $p(\overline{x})$ of $\pi(\overline{x})$ to BC such that $(p, C) \in I_P$.*

Proof. Let $p_1, ..., p_s$, $0 < s < \omega$, be an enumeration of the complete (quantifier-free) extension of π to BC. Thus, we must prove that $(p_j, C) \in I_P$ for some $j \in [s]$. Towards a contradiction, suppose that for each $j \in [s]$, there are a structure $\mathcal{M}_j \in K_{BC}$ and $BC \subseteq D_j \subseteq M_j$ such that $\Delta^{\mathcal{M}_j}[p_j, C](D_j)$ is infinite. By the amalgamation property, then, we may assume that there are a structure $\mathcal{M} \in K_{D_1} \cap .... \cap K_{D_s}$ and $D_1, ..., D_s \subseteq D \subseteq M$, such that $\Delta^{\mathcal{M}}[p_j, C](D)$ is infinite for each $j \in [s]$. Then, of course, $\Delta^{\mathcal{M}}[\pi, C](D)$ must infinite, contradicting the assumption that $(\pi, C) \in I_P$. □

Lemma 8.8 (Partial I_P-right-transitivity). *Let $\mathcal{M}_0 \in K$ and $B_1, B_2, C \subseteq M_0$, and let $\pi(\overline{x})$ be an admissible type over B_1B_2C such that $\pi{\restriction}B_1C$ is a complete quantifier-free type over B_1C in $\mathcal{M}_0$. Suppose $(\pi, C) \in I_P$. Then $(\pi, B_1C) \in I_P$ and $(\pi{\restriction}B_1C, C) \in I_P$.*

Proof. Given the assumptions, we know that $(\pi\restriction B_1C, C) \in I_P$ by the I_P-monotonicity property, which we proved above, so it only remains to show that $(\pi, B_1C) \in I_P$. Suppose $\mathcal{M} \in K_{B_1B_2C}$ and $B_1B_2C \subseteq D \subseteq M$. Let $\mathcal{N} \in K_D^{\pi} \setminus \Delta_0^{\mathcal{M}}[\pi, C](D)$, and let $\mathcal{N} \models \pi(\overline{a})$ such that $\overline{a} \downarrow_{B_1B_2C} D/\mathcal{N}$ and $\overline{a} \downarrow_C D/\mathcal{N}$. Since $B_1 \subseteq D$, by Local Base-monotonicity, we have $\overline{a} \downarrow_{B_1C} D/\mathcal{N}$, so $\overline{a}$ witnesses the fact that $\mathcal{N} \in K_D^{\pi} \setminus \Delta_0^{\mathcal{M}}[\pi, B_1C](D)$. Therefore,

$$K_D^{\pi}/{\cong_D} \setminus \Delta^{\mathcal{M}}[\pi, C](D) \;\subseteq\; K_D^{\pi}/{\cong_D} \setminus \Delta^{\mathcal{M}}[\pi, B_1C](D)$$

and therefore

$$\Delta^{\mathcal{M}}[\pi, B_1C](D) \;\subseteq\; \Delta^{\mathcal{M}}[\pi, C](D)$$

As the right-hand side is a finite set by hypothesis, the lemma follows. □

Lemma 8.9 (I_P-Preservation of algebraic dependence I)**.** *Suppose $\mathcal{M}_0 \in K$ and $A, C \subseteq M_0$. If $(qtp(A/AC; \mathcal{M}_0), C) \in I_P$, then $A \subseteq cl(C; \mathcal{N})$ for almost every $\mathcal{N} \in K_{AC}$*

Proof. Let $\pi = qtp(\overline{a}/AC; \mathcal{M}_0)$ where $\overline{a}$ is an enumeration of A. Suppose $\mathcal{M} \in K_{AC}$ and $AC \subseteq D \subseteq M$. Let $\mathcal{N} \in K_D^{\pi} = K_D$. As $\overline{a} \downarrow_{AC} D/\mathcal{N}$ (which holds trivially), then we may assume (*) that $\overline{a} \downarrow_C D/\mathcal{N}$, and in particular,

$$[A \amalg D \,|\, HDesc(\mathrm{cl}(C))]_{H_P^{\mathcal{N}}}$$

Since $A \subseteq AC \subseteq D$, it follows that $A \subseteq HDesc(\mathrm{cl}(C; \mathcal{N}))$, so by definition, we have $A \subseteq \mathrm{cl}(C; \mathcal{N})$. As (*) holds for almost every $\mathcal{N} \in K_{AC}$, the lemma follows. □

Lemma 8.10 (I_P-Preservation of algebraic dependence II)**.** *Suppose $\mathcal{M}_0 \in K$ and $B, C \subseteq M_0$, and let $\pi(\overline{x})$ be a complete quantifier-free type over BC. Assume $(\pi, C) \in I_P$. For any extension π' of π to $cl(B; \mathcal{M}_0) \cup C$, $(\pi', C) \in I_P$.*

As we conducted our arguments in "language" of I_P, we summarize the results of this section in the following theorem in the "language" of $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}$.

Theorem 8.11. *$\mathop{\smile\!\!\!\!\!\!|}^{d}$ satisfies the following properties of an independence relation:*

1. *Invariance: $(A, B, C) \equiv^{qf} (A_1, B_1, C_1) \wedge A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B \Rightarrow A_1 \mathop{\smile\!\!\!\!\!\!|}^{d}_{C_1} B_1$.*
2. *Existence: $A \mathop{\smile\!\!\!\!\!\!|}^{d}_{B} B$.*
3. *Monotonicity: $A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B \wedge B_0 \subseteq B \Rightarrow A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B_0$.*
4. *Base-monotonicity: $A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B \wedge B_0 \subseteq B \Rightarrow A \mathop{\smile\!\!\!\!\!\!|}^{d}_{B_0C} B$.*
5. *Partial right-transitivity: $A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B_1B_2 \Rightarrow A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B_1 \wedge A \mathop{\smile\!\!\!\!\!\!|}^{d}_{B_1C} B_2$.*
6. *Preservation of algebraic dependence I: $A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} A \Rightarrow A \subseteq cl(C)$.*

7. *Preservation of algebraic dependence II:* $A \mathop{\smile\!\!\!\!\!\!\downarrow}^{d}_{C} B \Rightarrow A \mathop{\smile\!\!\!\!\!\!\downarrow}^{d}_{C} cl(B)$.

8. *Preservation of algebraic dependence III:* $A \subseteq cl(C) \Rightarrow A \mathop{\smile\!\!\!\!\!\!\downarrow}^{d}_{C} B$

9. *Extension:* $A \mathop{\smile\!\!\!\!\!\!\downarrow}^{d}_{C} B \wedge BC \subseteq B_1 \Rightarrow \exists A' \equiv^{qf}_{BC} A.\, A' \mathop{\smile\!\!\!\!\!\!\downarrow}^{d}_{C} B_1$.

8.4 Symmetry and full transitivity

Unsurprisingly, symmetry and full transitivity, even for I_P, do not seem to hold without some additional hypothesis. In this section, we define and analyze one such assumption: *Bounded unfolding-degrees with respect to* P (abbreviated BUDs$/P$). In later sections, we show that assuming the complexity of P over K is polynomially bounded, then K does indeed have BUDs$/P$.

Towards defining BUDs$/P$, we first define our notion of unfolding-degrees. Let $\mathcal{M}_0 \in K$, $B, E \subseteq M_0$ and $0 < n < \omega$; then, for any $d < \omega$, $\deg^n_P(B/E; \mathcal{M}_0) \geq d$ if there are $\mathcal{M} \in K_{BE}$ and $\overline{a}_1, ..., \overline{a}_d \in M^n$ such that

- $(\overline{a}_1, ..., \overline{a}_d)$ is a BE-indiscernible sequence (with respect to quantifier-free formulas);
- $[\overline{a}_i \amalg \overline{a}_j \,|\, HDesc(BE)]_{H^{\mathcal{M}}_P}$ whenever $1 \leq i < j \leq d$;
- $[\overline{a}_i \not\amalg B \,|\, HDesc(E)]_{H^{\mathcal{M}}_P}$ for all $i = 1, ..., d$.

We say that K has *bounded unfolding-degrees with respect to* P if there is a function $f_{\deg} : \mathbb{N}^3 \to \mathbb{N}$ such that

$$\deg^n_P(B/E; \mathcal{M}_0) \leq f_{\deg}(n, |BE|, |E|)$$

whenever $\mathcal{M}_0 \in K$, $B, E \subseteq M_0$ and $n \in \mathbb{N}$.

Lemma 8.12. *Let $\mathcal{M}_0 \in K$, $C \subseteq M_0$ and $\overline{a}, \overline{b} \in M_0^{<\omega}$, and let $p(\overline{x}, \overline{y}) = qtp(\overline{a}, \overline{b}/C; \mathcal{M}_0)$. Assume $\overline{a} \mathop{\smile\!\!\!\!\!\!\downarrow}^{d}_{C} \overline{b}$. If $\overline{b} \not\mathop{\smile\!\!\!\!\!\!\downarrow}^{d}_{C} \overline{a}$ and $\overline{b} \in cl(C\overline{a})$, then K does not have BUDs$/P$.*

Proof. Assuming $\overline{b} \not\mathop{\smile\!\!\!\!\!\!\downarrow}^{\mathrm{d}}_{C} \overline{a}$, let $\mathcal{M} \in K_{C\overline{a}}$ and $C \subseteq D \subseteq M \setminus rng(\overline{a})$ be such that

$$\Delta = \Delta^{\mathcal{M}}[p(\overline{a}, \overline{y}), C](D\overline{a})$$

is infinite. By the property of Preservation of algebraic dependence III, $\overline{b} \in \mathrm{cl}(C\overline{a})$ implies that $\overline{b} \mathop{\smile\!\!\!\!\!\!\downarrow}^{\mathrm{d}}_{C\overline{a}} D\overline{a}$. Let $\hat{\mathcal{N}} \in \Delta$ and $\overline{b}' \in p(\overline{a}, \hat{N})$. If $\mathrm{cl}(C\overline{a}; \hat{\mathcal{N}}) \subseteq E \subseteq \mathrm{cl}(D\overline{a}; \hat{\mathcal{N}})$, then $[\overline{b}' \amalg D\overline{a} \,|\, HDesc(\mathrm{cl}(E))]_{H^{\hat{\mathcal{N}}}_P}$ just because $\overline{b} \subseteq HDesc(\mathrm{cl}(E))$; hence, $\overline{b}' \downarrow_{C\overline{a}} D\overline{a}/\hat{\mathcal{N}}$, so $\overline{b}' \not\downarrow_{C} D\overline{a}/\hat{\mathcal{N}}$. (Note that this is more or less the proof of Preservation of algebraic dependence III.) Furthermore, if $\overline{a} \vDash qtp(\overline{a}/D\overline{b}; \hat{\mathcal{N}})$ in some $\hat{\mathcal{N}} \in K_{D\overline{b}}$, then $\overline{b} \downarrow_{C\overline{a}'} D\overline{a}'/\hat{\mathcal{N}}$ and $\overline{b} \not\downarrow_{C} D\overline{a}'/\hat{\mathcal{N}}$. Now, we consider an inductive construction:

Construction:

- Let $\mathcal{N}_0 = \mathcal{M}_0$ and $\overline{a}_0 = \overline{a}$.
- Suppose we are given $\mathcal{N}_0, ..., \mathcal{N}_i$ and $\overline{a}_0, ..., \overline{a}_i$ such that
  1. For each $j \leq i$, $\overline{a}_0, ..., \overline{a}_j \in p(N_j, \overline{b})$,
  $$qtp(\overline{a}_j/D\overline{b}; \mathcal{N}_j) = qtp(\overline{a}_j/D\overline{b}; \mathcal{N}_i) = qtp(\overline{a}_0/D\overline{b}; \mathcal{N}_0)$$
  2. For each $j < i$, $\overline{a}_{j+1} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{D\overline{b}} D\overline{b}\overline{a}_0...\overline{a}_j$
- Choose $\mathcal{N}_{i+1} \in K_{D\overline{b}\overline{a}_0...\overline{a}_i}$ and $\overline{a}_{i+1} \in p(N_{i+1}, \overline{b})$ such that
  $$qtp(\overline{a}_{i+1}/D\overline{b}\overline{a}_{<i}; \mathcal{N}_{i+1}) = qtp(\overline{a}_i/D\overline{b}\overline{a}_{<i}; \mathcal{N}_i)$$
  and $\overline{a}_{i+1} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{D\overline{b}} D\overline{b}\overline{a}_0...\overline{a}_i$.

By methods we've used several times before, we may assume that $(\mathcal{N}_i, \overline{a}_i)$ is coherent and indiscernible over $\mathrm{cl}(D\overline{b}; \mathcal{N}_0)$ and that $\mathrm{cl}(D\overline{b}; \mathcal{N}_i) = \mathrm{cl}(D\overline{b}; \mathcal{N}_0)$ for all $i < \omega$. By Monotonicity, we know that $\overline{a}_i \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{D\overline{b}} \overline{a}_j$ whenever $j < i < \omega$, so $[\overline{a}_i \amalg \overline{a}_j \,|\, HDesc(\mathrm{cl}(D\overline{b}))]_{H_P^{\mathcal{N}_k}}$ provided $i < j << k < \omega$. On the other hand, by Ramsey's theorem, there is an infinite subset $X \subseteq \omega$ and some $acl(C) \subseteq E \subseteq acl(D)$ such that $[\overline{a}_i \not\amalg \overline{b} \,|\, HDesc(\mathrm{cl}(E))]_{H_P^{\mathcal{N}_k}}$ whenever $i, k \in X$ and $i << k$. This suffices, then, to demonstrate that K does not have BUDs/P. $\square$

Lemma 8.13. *Assume K has BUDs/P. Let $\mathcal{M}_0 \in K$, $C \subseteq M_0$ and $\overline{a}, \overline{b} \in M_0^{<\omega}$, and let $p(\overline{x}, \overline{b}) = qtp(\overline{a}/C\overline{b}; \mathcal{M}_0)$. Let $K_{C\overline{b}}$ and $C \subseteq D \subseteq M \backslash rng(\overline{b})$, and suppose $\Delta^{\mathcal{M}}[p(\overline{x}, \overline{b}), C](D\overline{b})$ is infinite. Let $q(\overline{x}, \overline{b})$ be a complete extension of $p(\overline{x}, \overline{b})$ to $D\overline{b}$. Then, if $(q(\overline{x}, \overline{b}), C\overline{b}) \in I_P$, then $q(\overline{x}, \overline{b})$ is algebraic.*

Proof. Assume $\mathcal{N} \in \Delta = \Delta^{\mathcal{M}}[p(\overline{x}, \overline{b}), C](D\overline{b})$ and $\overline{a}' \in p(N, \overline{b})$ such that the type $q(\overline{x}, \overline{b})$ $= qtp(\overline{a}'/D\overline{b}; \mathcal{N})$ does not deviate over $C\overline{b}$.

Let $0 < n < \omega$ be given. Choose $n_1 < \omega$ large enough to ensure that if $\mathcal{N}' \in K_{D\overline{b}}$ and $\overline{a}_1, ..., \overline{a}_{n_1} \in q(N', \overline{b})$, then there are $\mathrm{cl}(C; \mathcal{N}') \subseteq E \subseteq \mathrm{cl}(D; \mathcal{N}')$ and $i_1 < \cdots < i_n \leq n_1$ such that $(\overline{a}_{i_1}, ..., \overline{a}_{i_n})$ is $\mathrm{cl}(D\overline{b}; \mathcal{N}')$-indiscernible and $[\overline{a}_{i_j} \not\amalg D\overline{b} \,|\, HDesc(\mathrm{cl}(E; \mathcal{N}'))]_{H_P^{\mathcal{N}'}}$ for all $j = 1, ..., n$.

Now, by hypothesis, we have $\overline{a} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{C\overline{b}} D\overline{b}$, and by Extension – and using the inductive construction of lemma 8.12 – we may choose a model $\mathcal{M}^* \in K_{D\overline{b}}$ and $\overline{a}_1, ..., \overline{a}_{n_1} \in q(M^*, \overline{b})$ such that $\overline{a}_{i+1} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{C\overline{b}} D\overline{a}_1...\overline{a}_i$ whenever $1 \leq i < n_1$. By our choice of n_1, we may recover $i_1 < \cdots < i_n \leq n_1$ and $\mathrm{cl}(C; \mathcal{M}^*) \subseteq E \subseteq \mathrm{cl}(D; \mathcal{M}^*)$ such that

$$[\overline{a}_{i_j} \amalg \overline{a}_{i_k} \,|\, HDesc(\mathrm{cl}(E\overline{b})]_{H_P^{\mathcal{M}^*}}$$

whenever $1 \leq j \leq k \leq n$, and

$$\overline{a}_{i_j} \not\amalg D\overline{a} \,|\, HDesc(\mathrm{cl}(E))]_{H_P^{\mathcal{M}^*}}$$

whenever $1 \leq j \leq n$.

Noting that n was given arbitrarily, it's straightforward to show from this that K cannot have BUDs/P, and this proves the lemma. $\square$

Lemma 8.14. *Let $\mathcal{M} \in K$ and $C \subseteq D \subseteq M$, and let $p(\overline{x})$ be a complete type over $(C; \mathcal{M})$. If every complete extension of $p(\overline{x})$ to D is algebraic, then $p(\overline{x})$ itself is algebraic.*

Proof. An easy application of the amalgamation property. $\square$

Theorem 8.15. *If K has BUDs/P, then $\mathop{\smile\!\!\!\!\!\!|}^{d}$ is symmetric: $A \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B \Rightarrow B \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} A$. More precisely, suppose $\mathcal{M}_0 \in K$, $C \subseteq M_0$ and $p(\overline{x}, \overline{y}) = qtp(\overline{a}\overline{b}/C; \mathcal{M}_0)$ for some $\overline{a} \in M_0^m$ and $\overline{b} \in M_0^n$; then, if $(p(\overline{x}, \overline{b}), C) \in I_P$, then $(p(\overline{a}, \overline{y}), C) \in I_P$.*

Proof. Towards a contradiction, suppose $\mathcal{M} \in K_{C\overline{a}}$ and $C \subseteq D \subseteq M \setminus rng(\overline{a})$ are such that

$$\Delta^{\mathcal{M}}[p(\overline{a}, \overline{y}), C](D\overline{a})$$

is infinite. By 8.13, we may assume that $p(\overline{a}, \overline{y})$ is non-algebraic, and by the property of Preservation of algebraic dependence III, we may select non-algebraic extension $q(\overline{a}, \overline{y})$ of $p(\overline{a}, \overline{y})$ to $D\overline{a}$. By 8.12, we may assume that $q(\overline{a}, \overline{y})$ deviates over $C\overline{a}$. That is, if $\overline{b}' \vDash q(\overline{a}, \overline{y})$, then $\overline{b}' \not\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{C\overline{a}} D\overline{a}$. Consequently, there is a model $\mathcal{M}' \in K_{D\overline{a}}$ and a subset $D \subsetneqq D' \subseteq M' \setminus rng(\overline{a})$ such that

$$\Delta^{\mathcal{M}'}[q(\overline{a}, \overline{y}), C\overline{a}](D'\overline{a})$$

is infinite. Following the construction form 8.12, we obtain an infinite sequence $(\mathcal{N}_i, \overline{b}_i)_{i<\omega}$, coherence and indiscernible over $D\overline{a}'$ such that

$$\overline{b}_i \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{D'\overline{a}} D'\overline{a}\overline{b}_{<i}$$

whenever $i < \omega$, but

$$\overline{b}_i \downarrow_{D\overline{a}} D'\overline{a}/\mathcal{N}_j \;\Rightarrow\; \overline{b}_i \not\downarrow_{C\overline{a}} D'\overline{a}/\mathcal{N}_j$$

whenever $i << j < \omega$. By Ramsey's theorem, there is an infinite subset $X \subseteq \binom{\omega}{2}$ such that either $\overline{b}_i \not\downarrow_{D\overline{a}} D'\overline{a}/\mathcal{N}_j$ for all $\{i < j\} \in X$, or $\overline{b}_i \not\downarrow_{C\overline{a}} D'\overline{a}/\mathcal{N}_j$ for all $\{i < j\} \in X$. In either case, we derive a contradiction to BUDs/P. $\square$

Lemma 8.16 (Partial left-transitivity (BUDs/P))**.** *Let $\mathcal{M}_0 \in K$, $B, C \subseteq M_0$, and $\overline{a}_1, \overline{a}_2 \in M_0^{<\omega}$. If $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B$ and $\overline{a}_2 \mathop{\smile\!\!\!\!\!\!|}^{d}_{\overline{a}_1 C} B$, then $\overline{a}_1\overline{a}_2 \mathop{\smile\!\!\!\!\!\!|}^{d}_{C} B$.*

Proof. Towards a contradiction, suppose $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{C} B$, $\overline{a}_2 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{\overline{a}_1 C} B$, but $\overline{a}_1\overline{a}_2 \not\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{C} B$.

Claim. *We may assume that $\Delta^{\mathcal{M}_0}[qtp(\overline{a}_1\overline{a}_2/BC), C](BC)$ is infinite.*

proof of claim. Let $\mathcal{M} \in K_{BC}$ and $BC \subseteq D \subseteq M$ such that

$$\Delta^{\mathcal{M}}[qtp(\overline{a}_1\overline{a}_2/BC, C](D)$$

is infinite. Let $E = D \setminus BC$. By Extension, we may assume that $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_C BE$, and we can find $\overline{a}_2'$ (in some $\mathcal{M}' \in K_{D\overline{a}_1}$) such that $\overline{a}_2' \equiv^{\mathrm{qf}}_{BC\overline{a}_1} \overline{a}_2$. Up to renaming elements, we may then assume $\overline{a}_1 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_C BE$ and $\overline{a}_2 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{\overline{a}_1 C} BE$, but

$$\Delta^{\mathcal{M}'}[qtp(\overline{a}_1\overline{a}_2/BCE), C](BCE)$$

is infinite. □

By the property of Preservation of algebraic dependence III, we know that $qtp(\overline{a}_1\overline{a}_2/C)$ is non-algebraic, and it follows that $qtp(\overline{a}_1/C)$ or $qtp(\overline{a}_2/C\overline{a}_1)$ is non-algebraic (possibly both).

Claim. *$qtp(\overline{a}_1\overline{a}_2/BC)$ is non-algebraic.*

proof of claim. Towards a contradiction, suppose $qtp(\overline{a}_1\overline{a}_2/BC)$ is algebraic. Then the type $qtp(\overline{a}_1/BC)$ is algebraic (applying the amalgamation property), and also $qtp(\overline{a}_2/BC\overline{a}_1)$ is algebraic. If $qtp(\overline{a}_1/C)$ is non-algebraic, then $\overline{a}_1 \not\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_C B$; and if $qtp(\overline{a}_2/C\overline{a}_1)$ is non-algebraic, then $\overline{a}_2 \not\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{\overline{a}_1 C} B$. In either case, we have a contradiction to the hypothesis of the lemma. □

From the second claim, it is, then, straightforward to derive a contradiction to BUDs/P. □

Theorem 8.17 (BUDs/P). *Let $\mathcal{M}_0 \in K$ and $A, B_1, B_2, C \subseteq M_0$. Then, $A \mathop{\smile\!\!\!\!\!\!|}^{d}_C B_1$ and $A \mathop{\smile\!\!\!\!\!\!|}^{d}_{B_1 C} B_2$ if and only if $A \mathop{\smile\!\!\!\!\!\!|}^{d}_C B_1 B_2$.*

Proof. The proof is a simple derivation:

$$\begin{aligned} A \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_C B_1 \wedge A \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{B_1 C} B_2 &\Rightarrow B_1 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_C A \wedge B_2 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_{B_1 C} A \\ &\Rightarrow B_1 B_2 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_C A \\ &\Rightarrow A \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{d}}_C B_1 B_2 \end{aligned}$$

by Symmetry followed by Partial left-transitivity followed by Symmetry again. □

One last intermediate observation of this section (given without the now-routine proof) is the following:

Lemma 8.18 (BUDs/P). *Let $\mathcal{M}_0 \in K$ and $A, C \subseteq M_0$. Suppose $(\mathcal{M}_i, \overline{b}_i)_{i<\omega}$ is an AC-indiscernible coherent sequence. Then there is a function $g : \omega \to \omega$ such that $i \leq g(i)$ and, with $B_i = \bigcup \{rng(\overline{b}_j) : j < i\}$*

$$\overline{b}_{g(i)} \mathop{\smile\!\!\!\!\!\!|}^{d}_{CB_i} A$$

for all $i < \omega$.

The culmination of this section, then, is the following theorem, which one now derives easily from lemma 8.18:

Theorem 8.19. *Let P be an essentially inflationary* `invent`*-program over K, and suppose K has BUDs/P. Then $\mathop{\smile\!\!\!\!\!\!|}^{d}$ is an independence relation with local character. In particular, K is rosy.*

8.5 Boundedness properties of efficient programs

The reader may have observed that the notion of independence defined above does not have anything to do with the model-building problem explicitly. In this section, we finally return to the notion of a coherent solution of the model-building problem, and we show that an efficient coherent solution yields a local boundedness property which then implies BUDs/P.

Towards defining this local boundedness condition, we first consider a structure $\mathcal{M} \in K$ and $A \subseteq M$; we define

$$\texttt{res}_s(A;\mathcal{M}) = \left\{i < \omega : V_{i+1}(H_P^{\mathcal{M}}) \cap HDesc(A;\mathcal{M}) \ncong_A V_i(H_P^{\mathcal{M}}) \cap HDesc(A;\mathcal{M})\right\}$$

$$\texttt{res}_0(A;\mathcal{M}) = |\texttt{res}_s(A;\mathcal{M})|$$

The set $\texttt{res}_s(A;\mathcal{M})$ and the number $\texttt{res}_0(A;\mathcal{M})$ must exist just because $\texttt{resp}(P,\mathcal{M})$ is defined by assumption, and it's easy to see that if

$$(\texttt{res}_0(A;\mathcal{M}), v) \in HDesc(A;\mathcal{M})$$

then v is of the form $(0, \overline{a}, X, +)$. This number should, presumably, be called the *resolution rank* of A in $\mathcal{M}$ with respect to P. Next, we define

$$\texttt{res}(A;\mathcal{M}) = \begin{cases} \max\{\texttt{res}_0(A;\mathcal{N}) : \mathcal{N} \in K_A\} & \text{if the maximum exists} \\ \infty & \text{otherwise} \end{cases}$$

and call this number the resolution rank of the diagram $qtp(A;\mathcal{M})$. (It is easy to see that `res` is an invariant of the isomorphism type $qtp(A;\mathcal{M})$.) Finally, we say that K *is locally bounded with respect to P* just in case there is a function $f_{\texttt{res}} : \mathbb{N} \to \mathbb{N}$ such that

$$\texttt{res}(A;\mathcal{M}) \leq f_{\texttt{res}}(|A|)$$

whenever $\mathcal{M} \in K$ and $A \subseteq M$. The first fact involving this local boundedness property asserts an unsurprising connection with BUDs/P.

Observation. *Let $\mathcal{M} \in K$ and $A \subseteq M$. Suppose $j \in \texttt{res}_s(A;\mathcal{M})$, and let $h(j) = |\{i \in \texttt{res}_s(A;\mathcal{M}) : i < j\}|$. Then each $v \in V_j(H_P^{\mathcal{M}}) \cap HDesc(A;\mathcal{M})$ has no more than $\texttt{tower}_2(l_P, h(j))$ nontrivial predecessors (i.e. not in $HDesc(A;\mathcal{M})$), where $\texttt{tower}_2(x,1);= 2^x$ and $\texttt{tower}_2(x,n+1) := 2^{\texttt{tower}_2(x,n)}$.*

Proposition 8.20. *Suppose K is locally bounded with respect to P. Then K has BUDs/P.*

Proof. Towards a contradiction, suppose K does not have BUDs/P; in particular, let $\mathcal{M}_0 \in K$, $B, E \subseteq M_0$, and let $(\mathcal{M}_n, a_n)_{n<\omega}$ be a coherent BE-indiscernible sequence in K_{BE}. Let $r = f_{\mathtt{res}}(|E|)$. Observe that if $n < \omega$ and if v is in $V(H_P^{\mathcal{M}_n}) \setminus HDesc(A; \mathcal{M}_n)$ but has a descendent in $HDesc(A; \mathcal{M}_n)$, then v has a descendent in

$$HDesc_i(E; \mathcal{M}_n) = V_i(H_P^{\mathcal{M}_n}) \cap HDesc(A; \mathcal{M}_n)$$

for some $i \in \mathtt{res}_s(A; \mathcal{M}_n)$. By the pigeonhole principle, we may assume that for all $m, n < \omega$,

- $\mathtt{res}_0(A; \mathcal{M}_m) = \mathtt{res}_0(A; \mathcal{M}_n)$ – say

$$\mathtt{res}_0(A; \mathcal{M}_m) = \{i_1^m < \cdots < i_r^m\}$$
$$\mathtt{res}_0(A; \mathcal{M}_n) = \{i_1^n < \cdots < i_r^n\}$$

- For all $j = 1, ..., r$, $HDesc_{i_j^m}(A; \mathcal{M}_m) \cong_A HDesc_{i_j^n}(A; \mathcal{M}_n)$ (up to renumbering levels in the obvious way).

We may, therefore, think of all of the unions $\bigcup_{j=1}^r HDesc_{i_j^n}(E, \mathcal{M}_n)$ as a single finite set $\hat{E}$ which is constant relative to n.

For each $n < \omega$ and $i \leq n$, let τ_i^n be a trail from a_i to B in $H_P^{\mathcal{M}_n}$ which is not blocked by $HDesc(A; \mathcal{M}_n)$. Note every head-to-head vertex of τ_i^n *must* have a descendent in $HDesc(E; \mathcal{M}_n)$ (hence in $\hat{E}$), and τ_i^n must have at least one head-to-head vertex because its endpoints are on level zero while $H_P^{\mathcal{M}_n}$ has edges only between levels.

Now, for each $n < \omega$, define $g_n : [n] \to \hat{E}$ so that $g_n(i)$ is the youngest descendant of the first head-to-head vertex of τ_i^n (reading from a_i towards B). For every $s < \omega$, there is an $n_s < \omega$ and an $e \in \hat{E}$ such that $|g_n^{-1}(e)| \geq s$. Thus, there is an $e_0 \in \hat{E}$ such that $|g_{n_s}^{-1}(e_0)| \geq s$ for infinitely many $s < \omega$, and therefore e_0 has unboundedly many nontrivial predecessors, contradicting the fact that e_0 can have at most $\mathtt{tower}_2(l_P, r+1)$ nontrivial predecessors. □

The next theorem, 8.23, is in some sense the "fundamental" theorem in showing that if K admits an efficient coherent solution of its model-problem, then $\mathop{\smile\!\!\!\!\!|}^{\mathrm{d}}$ is a genuine independence relation. Before stating the theorem, we recall the definition of "coherent solution" for the reader's convenience. A *weakly coherent solution* of the model-building problem for K is an HL-transformation (`invent`-program) P of type $\rho \to \rho_1$, where $\rho \subseteq \rho_1$ which satisfies the following requirements:

1. $K_\forall \subseteq dom(P)$, and $\mathtt{resp}(P, A_\mathcal{M}) \restriction \rho \in K$ if $\mathcal{M} \in K$ and $A \subseteq M$.[2]

[2] Recall that if $\mathcal{B}$ is a ρ_1-structure – as must be $\mathtt{resp}(P, A_\mathcal{M})$ – then $\mathcal{B} \restriction \rho$ is understood, here, to be the ρ-structure with universe $B_0 = \bigcup_{R \in \rho} fld(R^\mathcal{B})$ and $R^{\mathcal{B} \restriction \rho} = R^\mathcal{B} \cap B_0^{ari(R)}$ for each $R \in \rho$.

2. Suppose $\mathcal{M} \in K$ and $A \subseteq B \subseteq M$. Then
$$A_{\mathcal{M}} \leq \texttt{resp}(P, A_{\mathcal{M}}) \leq \texttt{resp}(P, B_{\mathcal{M}})$$

3. Suppose $\mathcal{M}_0 \in K$ and $A \subseteq M_0$. Then, for any quantifier-free type $p(x_1, ..., x_k)$ over $(A; \texttt{resp}(P, A_{\mathcal{M}_0}))$, if
$$\texttt{resp}(P, A_{\mathcal{M}_0})\restriction\rho \leq \mathcal{M} \in K$$
and $p(\overline{x})$ is algebraic, then $q(M^k) \subseteq ||\texttt{resp}(P, A_{\mathcal{M}_0})||^k$.

Again, P is a fully coherent solution if it is essentially inflationary (which we have assumed throughout this chapter) and $\rho \subseteq pub(P)$. Now, the argument for the theorem is almost self-evident.

Lemma 8.21. *Assume P is a fully coherent solution for K. Let $\mathcal{M}_0 \in K$, $A \subseteq B \subseteq M_0$, and $\mathcal{M} \in K_B$. Suppose $b \in ||HDesc(A, B_{\mathcal{M}})||$, and let*
$$q(x) = qtp\,(b/||\texttt{resp}(P; A_{\mathcal{M}})\restriction\rho||; \texttt{resp}(P, B_{\mathcal{M}}))$$
Then $q(x)$ is algebraic in the sense of $\texttt{resp}(P, K_\forall)$.

(We note that $q(x)$ is a type in the language of $pub(P) \supseteq \rho$, and furthermore, b need not *be an element of the set $||\texttt{resp}(P; A_{\mathcal{M}})\restriction\rho||$.)*

Corollary 8.22. *There is a function $g : \mathbb{N} \to \mathbb{N}$ such that for all $\mathcal{M} \in K$ and $A \subseteq B \subseteq M$, $\#||HDesc(A; B_{\mathcal{M}})|| \leq g(|A|)$.*

Theorem 8.23. *Suppose P is a fully coherent solution of the model-building problem for K. Then K is locally bounded with respect to P, and consequently, K has bounded unfolding-degrees with respect to P.*

Proof. The theorem follows immediately from 8.22 and the fact that P is inflationary for relations in $pub(P)$. □

8.6 At last, the main event

Theorem 8.24. *Let $K = \boldsymbol{fin}[T^G]$, where T is a complete k-variable theory with infinitely many finite models up to isomorphism.*

I. If T is constructible, then K is rosy.

II. T is efficiently constructible if and only if K is super-rosy.

Proof. By theorems 8.19 and 8.23 of this chapter, if T is constructible – i.e., if K admits a fully coherent solution to its model-building problem – then K is rosy. We argued in chapter 7, that if K is a super-rosy Fraïssé class, then K admits a *polynomial-time* fully coherent solution to its model-building problem – that is, T is efficiently constructible. Finally, if T is efficiently constructible, then by I, K is rosy, and by theorem 4.12, K has small algebraicity. Thus, by 6.8 of chapter 6, K is super-rosy. □

Bibliography

[1] Serge Abiteboul and Victor Vianu. Computing with first-order logic. *Journal of Computer and System Sciences*, 50(2):309–335, 1995.

[2] Hans Adler. A geometric introduction to forking and thorn-forking. *Journal of Mathematical Logic*, to appear.

[3] John T. Baldwin and Olivier Lessmann. Amalgamation properties and finite models in L^n-theories. (Unpublished), 2000.

[4] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, 2006.

[5] Anselm Blumer, A. Ehrenfeucht, David Haussler, and Manfred K Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965, 1989.

[6] Steven Buechler. Coordinatization in superstable theories I: stationary types. *Transactions of the American Mathematical Society*, 288(1):101–114, 1985.

[7] Steven Buechler. Coordinatization in superstable theories II. *Transactions of the American Mathematical Society*, 307(1):411–417, 1988.

[8] Jan Van Den Bussche, Dirk Van Gucht, Marc Andries, and Marc Gyssens. On the completeness of object-creating database transformation languages. *Journal of the ACM (JACM)*, 44(2):272–319, 1997.

[9] Gregory Cherlin, Leo Harrington, and Alistair Lachlan. $\aleph_0$-categorical, $\aleph_0$-stable structures. *Annals of Pure and Applied Logic*, 28:103–135, 1985.

[10] Clifton Ealy and Alf Onshuus. Characterizing rosy theories. *Journal of Symbolic Logic*, 72(4):919–940, 2007.

[11] Heinz-Dieter Ebbinghaus and Joerg Flum. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer-Verlag, 2000.

[12] Martin Grohe. Large finite structures with few L^k-types. *Information and Computation*, 179(2):250–278, 1997. Special issue: LICS'97.

[13] Martine Grohe. *Logic and Automata – History and Perspectives*, chapter : Logic, graphs and algorithms. Amsterdam University Press, 2007.

[14] Paul Helman, Bernard M. E. Moret, and Henry D. Shapiro. An exact characterization of greedy structures. *SIAM Journal of Discrete Mathematics*, 6(2):274–283, 1993.

[15] Wilfrid Hodges. *Model Theory*, volume 42 of *Encyclopedia of mathematics and its Applications*. Cambridge University Press, 1993.

[16] Tappani Hyttinen. Forking in finite models. Unpublished manuscript.

[17] Tappani Hyttinen. On stability in finite models. *Archive for Mathematical Logic*, 39:89–102, 2000.

[18] Bernhard Korte and Lászlό Lovász. Mathematical structures underlying greedy algorithms. In F. Gecseg, editor, *Fundamentals of Computation Theory: Proceedings of the 1981 International FCT-Conference*, number 117 in Lecture Notes in Computer Science, pages 205–209. Springer-Verlag, 1981.

[19] Leonid Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. Springer-Verlag, 2004.

[20] Marko Djordjevic (now Vera Koponen). Finite variable logic, stability and finite models. *Journal of Symbolic Logic*, 66(2):837–858, 2001.

[21] Marko Djordjevic (now Vera Koponen). Finite satisfiability and $\aleph_0$-categorical structures with trivial dependence. *Journal of Symbolic Logic*, 71(3):810–830, 2006.

[22] Alf Onshuus. Properties and consequences of thorn-independence. *Journal of Symbolic Logic*, 71(1):1–21, 2006.

[23] Alf Onshuus and Alex Usvyatsov. Orthogonality and domination in unstable theories. *arXiv:0909.3753*, 2009.

[24] Martin Otto. *Bounded Variable Logics and Counting: A Study in Finite Models*. Lecture Notes in Logic. Springer-Verlag, 1997.

[25] Christos Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, 1982.

[26] Judea Pearl and Thomas Verma. The logic of representing dependencies by directed graphs. Technical Report CSD 870004, R-79-II, University of California, Los Angeles, Cognitive Systems Laboratory, 1987.

[27] Saharon Shelah. *Classification Theory and the Number of Nonisomorphic Models*, volume 92 of *Studies in Logic and the Foundations of Mathematics*. North–Holland Publishing Co., Amsterdam, 2nd edition, 1990.

[28] Frank O. Wagner. *Simple Theories*, volume 503. Kluwer Academic Publishers, 2000.