Independence, Amalgamation, and Trees

by

Samuel Nicholas Ramsey

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Logic and the Methodology of Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Thomas Scanlon, Chair
Professor Jacob Gaboury
Professor Leo Harrington
Professor Pierre Simon
Professor John Steel

Spring 2018

Independence, Amalgamation, and Trees

Abstract

Independence, Amalgamation, and Trees

by

Samuel Nicholas Ramsey

Doctor of Philosophy in Logic and the Methodology of Science

University of California, Berkeley

Professor Thomas Scanlon, Chair

This thesis is concerned with developing a theory of model-theoretic tree properties. These properties are combinatorial properties of a formula or family or formulas that place strong constraints on the behavior of forking and dividing yet are compatible with certain forms of model-theoretic randomness. The most significant and intensively studied is *the tree property*, whose negation characterizes the simple theories, and a successful theory for simple theories was developed by Hrushovski, Kim, Pillay, and others, in the late 90s and early 2000s. Motivated by parallels with simplicity theory, we introduce a theory of independence called *Kim-independence* and present a structure theory for NSOP_1 theories in terms of it. This unifies and explains simplicity-like phenomena observed in several non-simple examples, such as existentially closed vector spaces with a bilinear form and ω-free PAC fields. This machinery also gives a streamlined method for establishing that a given theory is NSOP_1 and for showing that certain generic constructions preserve NSOP_1. We also develop techniques for the manipulation of tree indiscernibles to address several questions concerning the syntax of the related model-theoretic tree properties TP_1, weak k-TP_1, and the associated cardinal invariants.

In the first two chapters, we develop the syntax of model-theoretic tree properties. We consider two kinds of tree indiscernbles and establish several 'operations' on them that preserve indiscernibility. Using these tools, we prove in Chapter 1, joint with Artem Chernikov, that TP_1 is always witnessed by a formula in a single free variable, partially answering a question of Shelah, and that TP_1 is equivalent to weak k-TP_1, answering a question of Kim and Kim. In Chapters 1 and 2, we study cardinal invariants of a theory T called $\kappa_{\mathrm{cdt}}(T)$, $\kappa_{\mathrm{sct}}(T)$, and $\kappa_{\mathrm{inp}}(T)$, which measure approximations to TP, TP_1, and TP_2, respectively. We address two questions of Shelah concerning which relationships between the model-theoretic tree properties TP, TP_1, and TP_2 have quantitative analogues that are reflected in the values of these invariants. This is accomplished by a combination of model-theoretic techniques related to tree indiscernibles and set-theoretic tools from infinitary combinatorics.

Next, we focus on the tree property SOP_1 and develop an interpretation of this property in terms of independent amalgamation of types. This begins in Chapter 1, where we char-

acterize SOP_1 in terms of amalgamation and give a Kim-Pillay style criterion for NSOP_1, which is then applied to show many known examples are NSOP_1. This is supplemented and refined in Chapter 3, joint with Itay Kaplan, where we introduce Kim-independence, which generalizes non-forking independence in simple theories, and corresponds to non-forking at a generic scale. We show that in an NSOP_1 theory, Kim-independence is symmetric and satisfies the independence theorem and that, moreover, NSOP_1 is characterized by these properties. In Chapter 4, joint with Itay Kaplan and Saharon Shelah, we give a characterization of NSOP_1 in terms of the local character of Kim-independence. Then in Chapter 5, we prove that SOP_1 is witnessed by a formula in a single free variable.

Finally, we focus on applying the theory of independence for NSOP_1 theories to obtain new examples of NSOP_1 theories. Winkler showed that if T is a model-complete theory eliminating the quantifier $\exists^\infty$, then T has a model companion when viewed as a theory in a larger language, or when the language of T is enriched with Skolem functions, which may be viewed as the generic expansion of T by new relations, functions, and constants, or the generic Skolemization of T, respectively. In Chapter 6, we prove these constructions preserve NSOP_1 and characterize Kim-independence in the expansion in terms of Kim-independence in T. We then specialize our analysis to the model companion of the empty theory in an arbitrary language, arguably the simplest possible nonsimple NSOP_1 theory when the language contains a function symbol of arity at least 2, and give a complete description of forking and imaginaries in this theory. In Chapter 7, we relate model-theoretic tree properties of a PAC field to those of its absolute Galois group. Chatzidakis has shown that if the theory of the inverse system of the absolute Galois group of a PAC field is NSOP_n then so is the field, for $n \geq 3$. As NSOP_1, and NSOP_2 are defined very differently, extending this result to the $n = 1$ and $n = 2$ case requires a different argument, but we show that Chatzidakis's result extends to these cases by applying the theory of Kim-independence in the NSOP_1 case and analyzing strongly indiscernible trees in the NSOP_2 case.

To my mother and sister

Contents

List of Figures

Acknowledgments

Thanks to Tom Scanlon, for being such a source of stability through graduate school, supporting my work in myriad ways while giving me ample space to do my own thing. It was a privilege to be able to have so many generative conversations with Leo Harrington, who often seemed to come to an understanding my arguments before I myself had, and whose questions always put pressure on exactly what was going awry. I am thankful to John Steel for teaching wonderful set theory courses and for allowing me to talk about combinatorics in the inner model theory seminar. I am grateful for Pierre Simon for conversations which radically reshaped Chapter 4 of this thesis. Special thanks go to Jacob Gaboury for his willingness to be my outside member and to talk to me about the queer theory of computation.

I owe a profound debt to the logicans I worked with before arriving at Berkeley, Maryanthe Malliaris and Bill Tait in Chicago, Zoé Chatizdakis and Martin Hils in Paris, who have continued to play an important role in shaping and supporting my work. During the MSRI semester, I began working with Artem Chernikov, who treated me like a colleague long before he had any reason to do so, and this work laid the foundation for everything contained in this thesis. I am thankful for Byunghan Kim of course for inspiring Kim-independence and also for inviting me to Seoul and collaborating with me, as well to Jan Dobrowolski, Joon Kim, and Jung-Uk Lee for making my visit so enlightening and productive. In Jerusalem, I gained tremendously from conversations with Itay Kaplan and Saharon Shelah; to Itay I am especially grateful for inviting me and for being such a fun collaborator, to Saharon I am (obviously) especially grateful for the book.

At Berkeley, I am thankful for the logic community and innumerable conversations with my fellow students, especially David Casey, Ben Castle, Daniel Fremont, Gwyneth Harrison-Shermoen, Matthew Harrison-Trainor, Will Johnson, Daniel Lowengrub, James Moody, Noah Schweber, Benny Siskind, Lawrence Valby, Mariana Vicaria, James Walsh, Michael Wan, Kentarô Yamamoto, and Alex Youcis. Alex Kruckman in particular went above and beyond in meeting with me and commenting on my written work. I am also thankful for the past and current non-student members of the community, especially Jim Freitag, Wes Holliday, Silvain Rideau, and Carol Wood. Outside of Berkeley, I am very glad that I have Gabe Conant, Greg Cousins, Nadja Hempel, Lotte Kestner, Nadav Meir, Omer Mermelstein, Daoud Siniora, and Caroline Terry as friends and colleagues.

Lastly, I am thankful for the non-logician friends and family that were there and happy to talk about What It All Means. First and foremost, I want to acknowledge Jane Hu, who was proud of me when things went well, commiserated when they didn't, punctuated graduate school with lots of laughter, and inspired me with her lived example of what Sedgwick might have meant when she spoke of having a life where work and love are impossible to tell apart. It wouldn't have been fun without Forrest Brown, Kelsey Gee, Claire Hungerford, Janet Kwon, Mrinalini Ramesh, Josh Roselman, David Showalter, Emily Wang, and Brendan White. Lessons with Raghda Toman gave graduate school much needed balance and our conversations brighten every morning that we meet. Lastly, I am grateful to the Rhoades family and to my own family for, well, everything.

Part I

Introduction

Chapter 0

Introduction

0.1 Background

Much of mathematical logic is concerned with asking what kinds of mathematical things are the kinds of things we can classify. Figuring out what it means to classify something is part of the question. Shelah's *Classification Theory* [She90] begins by offering a precise and compelling suggestion for how to ask this question, and then proceeds to answer it. The foundational results of Morley and Baldwin-Lachlan showed how to take a logical hypothesis—$\aleph_1$-categoricity or strong minimality—and construct a theory of dimension from it, which in turn allowed one to determine the identity of a structure in terms of this single invariant. Shelah took this as a paradigm and proposed interpreting the classification problem with 'mathematical things' understood as models of a first-order theory and 'classification' interpreted as the assigment of cardinal invariants. And after nearly 600 pages, his book concludes by more or less completely characterizing when this is possible. It is an astonishing human achievement.

From Shelah's *Classification Theory* there emerged a general philosophy that, in classifying first-order theories, one should look for *dividing lines.* These are properties of theories that split the first-order theories into 'tame' and 'wild'. From this perspective, one should be able to prove interesting theorems about theories lying on either side—*structure theorems* for tame theories, *nonstructure theorems* for wild ones. Additionally, one hopes to find *inside* criteria for the dividing line, which characterize it in terms of syntax, and *outside* criteria, that make reference only to semantics, that is, to the models. Despite the distance that model theory has traveled from Shelah's original motivation, the dividing lines philosophy still forms the basis of the vocabulary in which we describe what model theory is ultimately trying to do. Within stability theory, most dividing lines were isolated as the possible reasons that a theory could fail to be classifiable, but their interest extended beyond their relevance to the motivating classification problem. Consequently, it became interesting to look for dividing lines even among the unstable theories and *simplicity* was the first to receive extensive study.

Simplicity theory emerged out of the confluence of three lines of research within model theory. The first came from work of Shelah, who named the simple theories in [She80] and gave a preliminary analysis of independence in simple theories via the notion of weak dividing. Although the tree property had been isolated by Shelah in [She90] in the course of his development of forking independence in stable theories, this paper was the first place in which theories without the tree property were studied in their own right. Shelah's aim here was to show that simplicity is a dividing line by giving an 'outside' set-theoretic characterization of the class of simple theories. He showed that, in a simple theory, the Boolean algebra of non-weakly-dividing formulas satisfies a certain chain condition. He then showed that, by forcing with this poset, one can construct very saturated models of a simple theory. More precisely, Shelah defined the *saturation spectrum* of a theory T to be the set of pairs (λ, κ) where $\lambda \geq \kappa$ are cardinals and every model of T of size $\leq \lambda$ has a κ-saturated elementary extension of size λ. A standard unions-of-chains argument shows that if $\lambda = \lambda^{<\kappa}$, then (λ, κ) is in the saturation spectrum of *any* theory. Shelah's discovery was that, assuming a forcing axiom $\mathrm{Ax}_0\mu$, the simple theories may be characterized as those theories whose saturation spectra contain some (λ, κ) with $\lambda \neq \lambda^{<\kappa}$. Although the question of what the saturation spectra of first-order theories can be may appear rather arcane, Shelah's chain condition was later distilled into the S1 property of an ideal, a central notion in the work of Hrushovski on approximate subgroups [Hru12], and Hrushovski's letter to Shelah [Hru97], in which he demonstrates that different simple unstable theories can have distinct saturation spectra, opened up the study of generalized amalgamation, which continues to be a central topic within pure model theory.

Secondly, work of Hrushovski [Hru91] and Hrushovski-Pillay [HP94] on structures related to pseudo-finite fields led to many of the basic notions in simplicity theory. Pseudo-finite fields where characterized by Ax as those fields that are perfect, have Galois group $\hat{\mathbb{Z}}$, and satisfy an algebraic condition called being *pseudo-algebraically closed* (PAC). A field K is PAC if every absolutely irreducible variety defined over K has a K-rational point. When K is perfect, K may be viewed as a definably closed substructure of its algebraic closure K^{alg}. Hrushovski observed that the PAC condition can be reformulated model-theoretically as the requirement that every multiplicity 1 formula (in the sense of the strongly minimal theory $\mathrm{Th}(K^{\mathrm{alg}})$) with parameters coming from K is realized in K. This condition makes sense and is elementary in an arbitrary strongly minimal theory with weak elimination of imaginaries and the definable multiplicity property, so Hrushovski set out to analyze PAC substructures of strongly minimal theories satisfying these additional technical conditions, together with a boundedness condition on the model-theoretic Galois group. Working at this level of abstraction required the analysis to center on the fundamental properties of dimension and genericity that the PAC substructure inherits from the ambient strongly minimal one. The work of Hrushovski and Hrushovski-Pillay proved many approximations to fundamental results of stability and stable group theory in an unstable setting by working locally and relating properties of the unstable structure to those of an ambient stable one.

The third line of model-theoretic research that led to simplicity theory came from Lachlan's program of classifying smoothly approximable structures. The starting point for this theory was the pioneering work on ω-stable $\aleph_0$-categorical theories done by Cherlin, Harrington, and Lachlan [CHL85]. This provided a point of contact between Lachlan's program of classifying homogeneous stable structures and Zilber's work on the geometric structure theory of $\aleph_1$-categorical theories. The Cherlin-Harrington-Lachlan analysis began with the classification of strictly minimal sets—that is, strongly minimal $\aleph_0$-categorical structures satisfying an additional technical hypothesis—and then, using the ambient dimension theory, to prove that an arbitrary ω-stable $\aleph_0$-categorical theory may be coordinatized in terms of them. One of the goals of this theory was to prove the non-finite axiomatizability of ω-stable $\aleph_0$-categorical structures and this was accomplished, in effect, by showing that they are smoothly approximable. This suggested that a similar structure theory might exist at the level of smoothly approximable structures. One significant obstacle to working at this level of generality, however, was the appearance of unstable structures. Kantor, Liebeck, and Macpherson [KLM] classified the primitive smoothly approximable structures and showed that they are all closely related to classical geometries over finite fields. These may be unstable: although an infinite dimensional vector space over $\mathbb{F}_p$ is ω-stable, an infinite dimensional vector space over $\mathbb{F}_p$ equipped with a nondegenerate bilinear form has the independence property. Replacements for the core notions of stability theory, such as stationarity and canonical bases, as well as finer notions from geometrical stability theory, like minimality and local modularity, had to be adapted to the setting of smoothly approximable structures in order to carry out a parallel theory in which an arbitrary smoothly approximable structure could be analyzed in terms of coordinatizing primitive geometries. This analysis was carried out by Cherlin and Hrushovski [CH03].

Much of the model-theoretic content in these developments was organized and explained by Kim and Pillay, who established a general theory of simple theories based upon the properties of the independence relation of non-forking [KP97] [Kim98] [Kim01]. They showed that a single combinatorial hypothesis—not having the tree property—was sufficient to recover almost all of the properties of non-forking independence observed in the special settings of bounded PAC substructures of strongly minimal sets and smoothly approximable structures. In particular, Kim proved Kim's lemma, which established that forking in a simple theory is always witnessed by a Morley sequence, from which the equality of forking and dividing and the symmetry of non-forking independence easily follow. Subsequently, Kim and Pillay proved the independence theorem and a characterization of non-forking independence in a simple theory in terms of abstract axioms of an independence relation. These breakthroughs clarified the fundamental phenomena at the heart of simplicity theory, allowing the theory to deepen. This led to the proliferation of new fundamental notions, such as hyperimaginaries, generalized amalgamation, stable forking, etc., which served as a driver of further model-theoretic research.

0.2 Problems from non-simple theories

So to recap the historical sketch given above, after Shelah's work, new examples and new phenomena created the need for generalizations of stability theory. These generalizations were subsequently systematized around properties of an abstract independence relation and these properties were shown to characterize and be characterized by simplicity. This point of view paved the way for the deepening of simplicity theory, which became a significant area of model-theoretic research in its own right, and served as the basis for further developments in the field. In what follows, we will explain how the theory presented in this thesis—the theory of Kim-independence and the theory of model-theoretic tree properties—may be motivated by a remarkably parallel situation in which previous work on non-simple examples can be unified and explained by a theory of independence, which in turn allows for new applications and suggests new structures to consider.

As in the case of the development of simple theories in the wake of stability, the new examples and new phenomena created that created need for new tools beyond simplicity were related to saturation, PAC structures, and classical geometries. The initial impetus to develop tree properties beyond simplicity comes from the $\triangleleft^*$-order introduced by Shelah to compare the the relative difficulty of producing saturated models of two theories. In their extensive study of this order [DS04], Džamonja and Shelah introduce the tree properties SOP_1 and SOP_2, as well as some variants of each, and study the relationship between these properties and maximality in the order $\triangleleft^*$. They observe that the theory of parametrized equivalence relations T^*_{feq} is neither simple nor maximal in $\triangleleft^*$, they show that maximality in $\triangleleft^*$ implies a property called SOP''_2, and also they observe that SOP_2 coincides with the property TP_1 introduced earlier by Shelah. Later, Shelah and Usvyatsov showed that SOP''_2 was equivalent to SOP_2 [SU08]. A complete description, modulo GCH, of the $\triangleleft^*$-maximal theories was given by Malliaris and Shelah who showed that they are exactly the SOP_2 theories [MS15a]. This was one of the many byproducts of their cofinality spectrum technology which also used to show SOP_2 theories are maximal in Keisler's order [MS16]. Their work on related problems has greatly clarified the relationship between model theoretic properties and the construction of saturated models, proving illuminating structure theorems among the simple theories and non-structure theorems for those with SOP_2. This points to an intriguing gap in our knowledge of the properties of nonsimple NSOP_2 theories living in-between.

A major source of new non-simple examples comes from unbounded PAC fields. Hrushovski's study of PAC substructures of strongly minimal sets imposed the hypothesis of *boundedness*. In the context of a PAC field K, this is equivalent to the requirement that the field K have finitely many degree n extensions for all n. Later, Chatzidakis showed that every unbounded PAC field is not simple, hence a PAC field is simple if and only if it is bounded [Cha99]. However, Chatzidakis developed a Galois-theoretic interpretation of non-forking independence in the context of ω-free PAC fields and, more generally, Frobenius fields and introduced a related independence relation called *weak independence* [Cha02]. She proved that Frobenius fields satisfy the independence theorem for weak independence. Moreover, this relation satisfies symmetry and extension, and therefore closely resembles non-forking independence in a simple theory, despite the fact that ω-free PAC fields and Frobenius fields are not simple.

Infinite-dimensional vector spaces over an algebraically closed field with a nondegenerate bilinear form are an another important example of an non-simple tame structure. These were given extensive study by Granger [Gra99], who proved quantifier-elimination in a reasonable language, showed that they are not simple, and developed a theory of independence for them. He defined Γ-independence in these structures and showed that it is symmetric, transitive, and stationary. Stationarity is a very strong form of the independence theorem, which can hold for non-forking independence only in a stable theory. Granger's examples represent one way of considering analogues of the classical geometries considered in the smoothly approximable case over an infinite field. Another way comes from work of Liebeck, Macpherson, and Tent who study and essentially classify primitive infinite ultraproducts of finite permutation groups, finding again a close connection with classical geometries [LMT09]. They observe in the conclusion of their paper that these ultraproducts should give rise to tame structures, but the appropriate notion of tameness and the associated structure theory are lacking:

> Initially, we hoped for a close connection between primitive ultraproducts (X^*, G^*) of finite permutation groups and simple theories, analogous to the smoothly approximable structures... One might have hoped that there is a supersimple structure M^* with domain X^* such that $G^* = \mathrm{Aut}(M^*)$, or, better (to avoid problems with field automorphisms), so that $\mathrm{Aut}(M^*) \leq G \leq N_{\mathrm{Sym}(X)}(\mathrm{Aut}(M))$. The latter seems correct, with the exception of cases where ultraproducts of unbounded L-rank symplectic, orthogonal or unitary groups, over unbounded fields, are involved. It was shown by Grainger [sic] that the theories of infinite-dimensional vector spaces carrying a non-degenerate sesquilinear form, over an infinite field, parsed in a two-sorted language, do not have simple theory. In Grainger's [sic] thesis some independence theory is developed for such structures (over an algebraically closed field), so there may be a reasonable model theory for all such structures M^*.

A model-theoretic treatment of classical geometries over infinite fields, then, seems to require an independence theory capable of generalizing and explaining the observations of Granger in the test case of a bilinear form over an algebraically closed field.

0.3 What we do

Syntax

A well-known theorem of stability theory says that an unstable theory is unstable for one of two conceptually distinct reasons. If T is unstable, then either T interprets a partial order with infinite chains, i.e. T has the *strict order property*, or T interprets a graph that contains the random graph as an induced subgraph, i.e. T has the *independence property*. In other words, if a theory has a lot of types, then they must either come from something very structured—cuts in an order—or something highly random—partitions of a random graph. Remarkably, this kind of dichotomy reappears at the level of forking. A sequence of nested intervals in a dense linear order may be taken as a paradigm for 'structured' forking, while a choice of equivalence classes in distinct equivalence relations in the theory T^*_{feq} of random parametrized equivalence relations may be regarded as a canonical example of 'random' forking. Shelah proves that theories with lots of forking or dividing give rise to a pattern of dividing that resembles one of these two examples. In order to state this precisely, we will need the definitions for Shelah's cardinal invariants related to the tree property:

Definition 0.3.1. 1. A *cdt-pattern of height* κ is a sequence of formulas $\varphi_i(x; y_i)$ ($i < \kappa, i$ successor) and numbers $n_i < \omega$, and a tree of tuples $(a_\eta)_{\eta\in\omega^{<\kappa}}$ for which

a) $p_\eta = \{\varphi_i(x; a_{\eta|i}) : i \text{ successor }, i < \kappa\}$ is consistent for $\eta \in \omega^\kappa$

b) $\{\varphi_i(x; a_{\eta\frown\langle\alpha\rangle}) : \alpha < \omega, i = l(\eta) + 1\}$ is n_i-inconsistent.

2. An *inp-pattern of height* κ is a sequence of formulas $\varphi_i(x; y_i)$ ($i < \kappa$), sequences $(a_{i,\alpha} : \alpha < \omega)$, and numbers $n_i < \omega$ so that

   a) for any $\eta \in \omega^\kappa$, $\{\varphi_i(x; a_{i,\eta(i)}) : i < \kappa\}$ is consistent.

   b) For any $i < \kappa$, $\{\varphi_i(x; a_{i,\alpha}) : \alpha < \omega\}$ is n_i-inconsistent.

3. An *sct-pattern of height* κ is a sequence of formulas $\varphi_i(x; y_i)$ ($i < \kappa$) and a tree of tuples $(a_\eta)_{\eta\in\omega^{<\kappa}}$ so that

   a) For every $\eta \in \omega^\kappa$, $\{\varphi_\alpha(x; a_{\eta|\alpha}) : 0 < \alpha < \kappa, \alpha \text{ successor}\}$ is consistent.

   b) If $\eta \in \omega^\alpha$, $\nu \in \omega^\beta$, α, β are successors, and $\nu \perp \eta$ then $\{\varphi_\alpha(x; a_\eta), \varphi_\beta(x; a_\nu)\}$ are inconsistent.

4. For $X \in \{\text{cdt}, \text{sct}, \text{inp}\}$, we define $\kappa^n_X(T)$ be the first cardinal κ so that there is no X-pattern of height κ in n free variables. We define $\kappa_X(T) = \sup\{\kappa^n_X\}$.

A cdt-pattern is a combinatorial configuration that emerges from instances of dividing. The sct-patterns and inp-patterns may be viewed as the extremal forms that a cdt-pattern can take. An sct-pattern is a cdt-pattern in which inconsistency is maximized and an inp-pattern is one in which inconsistency is minimized. Shelah proves that if $\kappa_{\text{cdt}}(T) = \infty$ then either $\kappa_{\text{sct}}(T) = \infty$ or $\kappa_{\text{inp}}(T) = \infty$ [She90]. By the pigeonhole principle and compactness, for $X \in \{\text{cdt}, \text{sct}, \text{inp}\}$, saying that $\kappa_X(T) = \infty$ is equivalent to saying that there is an infinite X-pattern in a single formula. The assertions $\kappa_{\text{cdt}}(T) = \infty$, $\kappa_{\text{sct}}(T) = \infty$, and $\kappa_{\text{inp}}(T) = \infty$ may be taken as the definition of T having the tree property (TP), the tree property of the first kind (TP_1), and the tree property of the second kind (TP_2), respectively.

One issue raised by Shelah in *Classification Theory* is the extent to which theorems about the tree properties TP, TP_1, and TP_2 have 'quantitative' analogues in terms of the cardinal invariants $\kappa_{\mathrm{cdt}}(T)$, $\kappa_{\mathrm{sct}}(T)$, and $\kappa_{\mathrm{inp}}(T)$. One particular question he asked in this vein is whether the equality $\kappa_{\mathrm{cdt}}(T) = \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$ for all T. In Chapter 1, joint with Artem Chernikov, we address this question in the case of countable theories. We introduce several 'operations' on indiscernible trees and show these preserve indiscernibility. Then using these operations, we study the way that a cdt-pattern of a given height can be manipulated to yield patterns with additional properties. In the case of a countable theory, the invariant $\kappa_{\mathrm{cdt}}(T)$ can only take on the values $\aleph_0, \aleph_1$, and ∞ and we answer the question positively for the countable theories with an argument tailored to the two cases not already covered by Shelah's theorem. In the finitary case $\kappa^n_{\mathrm{cdt}}(T) = \aleph_0$ where T has arbitrarily large finite cdt-patterns but no infinite one, we recover a quantitive form of Shelah's dichotomy. But we show that from an infinite cdt-pattern we can construct an infinite sct-pattern, hence $\kappa_{\mathrm{cdt}}(T) \geq \aleph_1$ if and only if $\kappa_{\mathrm{sct}}(T) \geq \aleph_1$. This has the curious consequence that the natural candidate for the definition of a 'super-NTP$_1$' theory—namely $\kappa_{\mathrm{sct}}(T) = \aleph_0$—coincides with supersimplicity. We also make use of the tree-indiscernible methods to prove that TP$_1$ is always witnessed by a formula in a single free variable, answering the local case of a another question of Shelah, who asked if $\kappa_{\mathrm{sct}}(T) = \kappa^1_{\mathrm{sct}}(T)$.

In Chapter 2, we consider the question of whether $\kappa_{\mathrm{cdt}}(T) = \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$ without any restriction on the cardinality of T. We construct a family of counterexamples, using *strong colorings* studied by Shelah in the context of pcf theory. From a strong coloring, we build a theory which can be obtained as a union of theories in finite reducts of the (uncountable) language, each of which is stable, $\aleph_0$-categorical, and has quantifier-elimination. In some sense, the gap between global complexity and local tameness is maximized in these examples so that the formulas in a cdt-pattern can be carefully analyzed and large sct- and inp-patterns can be ruled out. We also use a similar construction to answer a related technical question of Shelah concerning the impact of these cardinal invariants on saturation of ultrapowers.

Kim-independence

The theory of Kim-independence and the associated structure theory for NSOP$_1$ theories forms the heart of this thesis. To start, we give the definition of NSOP$_1$:

Definition 0.3.2. A formula $\varphi(x; y)$ has SOP$_1$ if there is a collection of tuples $(a_\eta)_{\eta \in 2^{<\omega}}$ satisfying the following:

- For all $\eta \in 2^\omega$, $\{\varphi(x; a_{\eta|n}) : n < \omega\}$ is consistent.
- For all $\eta, \nu \in 2^{<\omega}$, if $(\eta \wedge \nu) \frown \langle 0 \rangle \trianglelefteq \eta$ and $(\eta \wedge \nu) \frown \langle 1 \rangle = \nu$, then $\{\varphi(x; a_\eta), \varphi(x; a_\nu)\}$ is inconsistent.

A theory T has SOP$_1$ if some formula has SOP$_1$ modulo T. We say T is NSOP$_1$ if does not have SOP$_1$.

This rather bizarre-looking property was introduced by Džamonja and Shelah and later discussed by Shelah and Usvyatsov, but it remained unclear if there were any meaningful consequences to either having or lacking this SOP_1. We show that it may characterized in several distinct ways in terms of a natural independence relation and that $NSOP_1$ provides a worthwhile setting in which one can treat model-theoretically several structures of general mathematical interest.

Kim showed that many of the desirable properties of non-forking independence in a simple theory in fact *characterize* simplicity. A theory in which non-forking independence satisfies symmetry, (right) transitivity, Kim's lemma, or local character must be simple [Kim01]. Consequently, it is somewhat surprising that the property thought of as most characteristic of simplicity-like behavior, the *independence theorem*, does not imply simplicity. This was first explicitly observed by Chatzidakis, who showed that the independence theorem for non-forking independence (in its usual formulation as well as a dual version) holds in ω-free PAC fields, which were known to be non-simple by her earlier work. Despite the fact that the independence theorem does not imply simplicity, Kim observed, modulo the existence of a measurable cardinal, that the independence theorem for non-forking independence implies that the theory is NTP_1 [Kim01]. Later, Chernikov removed the set-theoretic hypothesis [Che14]. By a careful analysis of the proof, together with Chernikov, we noticed that from a weaker hypothesis we could obtain a stronger conclusion: if T has SOP_1 then the independence theorem fails in a strong way. Moreover, we showed this strong failure of the independence theorem must always come from an instance of SOP_1, yielding a new characterization of the $NSOP_1$ theories in terms of amalgamation of types. As a corollary, we obtained the following Kim-Pillay-style criterion for a theory being $NSOP_1$:

Proposition 0.3.3. *[Proposition 1.5.8] Assume there is an $Aut(\mathbb{M})$-invariant independence relation $\mathop{\smile\!\!\!\!\!\!|}$ on small subsets of the monster $\mathbb{M} \models T$ such that it satisfies the following properties, for an arbitrary $M \models T$ and arbitrary tuples from $\mathbb{M}$.*

1. *Strong finite character: if $a \not\mathop{\smile\!\!\!\!\!\!|}_M b$, then there is a formula $\varphi(x, b, m) \in tp(a/bM)$ such that for any $a' \models \varphi(x, b, m)$, $a' \not\mathop{\smile\!\!\!\!\!\!|}_M b$.*
2. *Existence over models: $M \models T$ implies $a \mathop{\smile\!\!\!\!\!\!|}_M M$ for any a.*
3. *Monotonicity: $aa' \mathop{\smile\!\!\!\!\!\!|}_M bb' \implies a \mathop{\smile\!\!\!\!\!\!|}_M b$.*
4. *Symmetry: $a \mathop{\smile\!\!\!\!\!\!|}_M b \iff b \mathop{\smile\!\!\!\!\!\!|}_M a$.*
5. *Independent amalgamation: $c_0 \mathop{\smile\!\!\!\!\!\!|}_M c_1$, $b_0 \mathop{\smile\!\!\!\!\!\!|}_M c_0$, $b_1 \mathop{\smile\!\!\!\!\!\!|}_M c_1$, $b_0 \equiv_M b_1$ implies there exists b with $b \equiv_{c_0 M} b_0$, $b \equiv_{c_1 M} b_1$.*

Then T is $NSOP_1$.

This criterion was then used to show that the theory of parametrized equivalence relations, infinite-dimensional vector spaces over an algebraically closed fields, and ω-free PAC fields all have an $NSOP_1$ theory. In each case, the notion of independence for these structures was *ad hoc*, defined in combinatorial or algebraic terms particular to the structures. In order to complete the analogy with the theorem of Kim and Pillay, however, it was necessary to develop a theory of independence in a general NSOP_1 theory, which would specialize to these *ad hoc* notions of independence in these particular theories.

This challenge is met by the theory of Kim-independence, a notion of independence defined in terms of dividing at a generic scale. The definition of Kim-independence emerged out of an attempt to better understand an argument of Malliaris and Shelah that characterized the NSOP_2 theories in terms of a notion they called a *higher formula* [MS15a]. Though this differs from their presentation, a higher formula may be described as a triple $(\varphi(x,y), A, q)$ where A is a set of parameters, q is a global A-finitely satisfiable type, and $\varphi(x;y)$ is a formula that does not divide with respect to any Morley sequence in q over A, or, to use Hrushovski's terminology, does not q-divide over A [Hru12]. We observed that NSOP_1 is equivalent to the statement that if q and r are global A-invariant types with $q|_A = r|_A$, then $(\varphi(x,y), A, q)$ is higher if and only if $(\varphi(x;y), A, r)$ is higher. In the argument, A-finite satisfiability was used merely to ensure A-invariance, so NSOP_1 is additionally equivalent to the statement that if q and r are A-invariant types with $q|_A = r|_A$ then, for any φ, $\varphi(x;y)$ q-divides over A if and only if $\varphi(x;y)$ r-divides over A. This can be viewed as a version of Kim's lemma for the following notion of dividing:

Definition 0.3.4. We say that a formula $\varphi(x;b)$ *Kim-divides* over A if there is some A-invariant global type $q \supseteq \mathrm{tp}(b/A)$ and Morley sequence $\langle b_i : i < \omega \rangle$ over A in q so that $\{\varphi(x;b_i) : i < \omega\}$ is inconsistent. The formula $\varphi(x;b)$ *Kim-forks* over A if $\varphi(x;b) \vdash \bigvee_{i<k} \psi_i(x;c^i)$ and each $\psi_i(x;c^i)$ Kim-divides over A. A type Kim-forks if it implies a formula which does. If $\mathrm{tp}(a/Ab)$ does not Kim-fork over A, we write $a \mathop{\smile\!\!\!\!\!\mid}^K_A b$.

In an NSOP_1 theory, Kim-dividing over a model M is dividing that is witnessed by some, equivalently every, Morley sequence in global M-invariant types over M. The interpretation of higher formulas and q-dividing in the NSOP_1 context in terms of the independence notion $\mathop{\smile\!\!\!\!\!\mid}^K$ was inspired by a suggestion of Kim for a notion of independence in the NTP_1 setting [Kim09]. This, together with the fact that $\mathop{\smile\!\!\!\!\!\mid}^K$ satisfies an analogue of Kim's lemma in NSOP_1 theories, is why the notion bears Kim's name.

In Chapter 3, joint with Itay Kaplan, we develop fairly comprehensive theory of Kim-independence in the NSOP_1 setting. The following is our main theorem, which gives several characterizations of NSOP_1 in terms of $\mathop{\smile\!\!\!\!\!\mid}^K$, which gives compelling evidence that SOP_1 is a dividing line:

Theorem 0.3.5. *[Theorem 3.8.1] The following are equivalent for the complete theory T:*

1. *T is $NSOP_1$*
2. *Kim's lemma for Kim-dividing: For every model $M \models T$ and $\varphi(x;b)$, if $\varphi(x;y)$ q-divides for some global M-invariant $q \supseteq tp(b/M)$, then $\varphi(x;y)$ q-divides for* every *global M-invariant $q \supseteq tp(b/M)$.*
3. *Symmetry over models: for every $M \models T$, then $a \mathop{\smile\!\!\!\!\!\mid}^K_M b$ if and only if $b \mathop{\smile\!\!\!\!\!\mid}^K_M a$.*

4. *Independence theorem over models: if $M \models T$, $a \equiv_M a'$, $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$, $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M c$, and $b \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M c$, then there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M bc$.*

Additionally, we refine the Kim-Pillay-style criterion for NSOP_1 into an abstract characterization of Kim-independence in an NSOP_1 theory:

Theorem 0.3.6. *[Theorem 3.9.1] Assume there is an $Aut(\mathbb{M})$-invariant ternary relation $\mathop{\smile\!\!\!\!\!\!\mid}$ on small subsets of the monster $\mathbb{M} \models T$ which satisfies the following properties, for an arbitrary $M \models T$ and arbitrary tuples from $\mathbb{M}$.*

1. *Strong finite character: if $a \not\mathop{\smile\!\!\!\!\!\!\mid}_M b$, then there is a formula $\varphi(x, b, m) \in tp(a/bM)$ such that for any $a' \models \varphi(x, b, m)$, $a' \not\mathop{\smile\!\!\!\!\!\!\mid}_M b$.*
2. *Existence over models: $M \models T$ implies $a \mathop{\smile\!\!\!\!\!\!\mid}_M M$ for any a.*
3. *Monotonicity: $aa' \mathop{\smile\!\!\!\!\!\!\mid}_M bb' \implies a \mathop{\smile\!\!\!\!\!\!\mid}_M b$.*
4. *Symmetry: $a \mathop{\smile\!\!\!\!\!\!\mid}_M b \iff b \mathop{\smile\!\!\!\!\!\!\mid}_M a$.*
5. *The independence theorem: $a \mathop{\smile\!\!\!\!\!\!\mid}_M b$, $a' \mathop{\smile\!\!\!\!\!\!\mid}_M c$, $b \mathop{\smile\!\!\!\!\!\!\mid}_M c$ and $a \equiv_M a'$ implies there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\!\!\!\!\!\!\mid}_M bc$*

Then T is $NSOP_1$ and $\mathop{\smile\!\!\!\!\!\!\mid}$ strengthens *$\mathop{\smile\!\!\!\!\!\!\mid}^K$—i.e. if $M \models T$, $a \mathop{\smile\!\!\!\!\!\!\mid}_M b$ then $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$. If, moreover, $\mathop{\smile\!\!\!\!\!\!\mid}$ satisfies*

6. *Witnessing: if $a \not\mathop{\smile\!\!\!\!\!\!\mid}_M b$ witnessed by $\varphi(x; b)$ and $(b_i)_{i<\omega}$ is a Morley sequence over M in a global M-invariant type extending $tp(b/M)$, then $\{\varphi(x; b_i) : i < \omega\}$ is inconsistent.*

then $\mathop{\smile\!\!\!\!\!\!\mid} = \mathop{\smile\!\!\!\!\!\!\mid}^K$.

This theorem serves as a tool for characterizing Kim-independence in concrete examples, even when finitely satisfiable and invariant types are not easy to directly understand. We apply this method to characterize Kim-independence in several theories, showing that Kim-independence coincides with the weak independence of Chatzidakis in Frobenius fields and is closely related to Granger's Γ-independence in infinite-dimensional vector spaces over algebraically closed fields with a nondegenerate bilinear form, which confirms the naturality of Kim-independence in this setting.

Several new technical notions were required for these results. Kim's lemma for Kim-dividing is a consequence of a synactic reformulation of SOP_1 in terms of an array of pairs, which is much easier to understand and manipulate than a tree. These syntactic reformulations were also essential for proving that SOP_1 is always witnessed in a single free variable, which is the content of Chapter 5. Familiar arguments from simplicity theory immediately adapt to show the equality Kim-forking and Kim-dividing in NSOP_1 theories, so $\mathop{\smile\!\!\!\!\!\!\mid}^K$ satisfies existence and extension over models. But symmetry of Kim-independence is proved after introducing tree Morley sequences which are indiscernible sequences sufficiently generic to witness Kim-dividing, but weaker than Morley sequences in invariant types and thus exist under more general hypotheses. The construction of tree Morley sequences is the most technically involved part of this thesis, as it requires the use of indiscernible trees and an inductive construction of certain ill-founded trees. These inductive constructions, and not just the existence of tree Morley sequences, are used repeatedly, particularly in the proof of the independence theorem for Kim-independence in NSOP_1 theories.

Chapter 4, joint with Itay Kaplan and Saharon Shelah, considers the local character of Kim-independence. In a simple theory, local character states that every type is a non-forking extension of a type of size $|T|$. Informally, this suggests that all the essential information about a realization of the type is contained in some small set of formulas. Despite the fact that Kim-independence gives a robust generalization of non-forking independence in simple theories, it lacks base monotonicity and this introduces subtleties into finding the right analogue of local character for Kim-independence. In a simple theory, if $p \in S(B)$ does not fork over some $A \subseteq B$, then p also does not fork over any set in the cone $\{C \subseteq B : A \subseteq B\}$. This fails in an NSOP_1 theory for Kim-forking, in general, but one would like to capture the idea that, in an NSOP_1 theory, a type should not Kim-fork over a 'large' number of small subsets. The appropriate notion of largeness is provided by the generalized club filter. Given a set X and a cardinal κ, a subset Z of $[X]^{\kappa}$ is called *club* if it is closed under $\subseteq$-increasing chains of length κ and for any $Y \in [X]^{\kappa}$ there is $Y' \in Z$ with $Y \subseteq Y'$. Accordingly, $Z \subseteq [X]^{\kappa}$ is *stationary* if it intersects every club subset. With these definitions in hand, we prove the following theorem:

Theorem 0.3.7 (Theorem 4.3.9)**.** *Suppose T is a complete theory. The following are equivalent:*

1. *T is $NSOP_1$.*

2. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a stationary subset of $[M]^{|T|}$.*

3. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a club subset of $[M]^{|T|}$.*

This gives yet another distinct characterization of NSOP_1 in terms of Kim-independence. We expect it will have applications in the future, in generalizing theorems from simplicity that make use of both local character and base monotonicity.

Examples

In the final section, we apply the theory developed earlier in the thesis to produce new examples of tame theories. In Chapter 6, joint with Alex Kruckman, we focus on 'generic structure' methods for producing new NSOP_1 theories from old ones. The precedent for these kinds of theorems was established by Chatzidakis and Pillay [CP98], who showed that adding a generic automorphism or generic predicate to a stable theory resulted in a simple theory, which was, in general, unstable. In the context of adding a generic predicate, it sufficed to assume the underlying theory is simple to show that adding a generic predicte produces a simple theory. In a different vein, Jeřábek studied the model companion of the empty theory in an arbitrary language, showing that is not simple when the language contains a function symbol of arity greater than 2 but is always NSOP_3. In an early preprint of [Jeř17], he showed that this theory is NSOP_3 for every choice of language and asked if it must always be NSOP_1. We address this question by situating it in a broader context, in which it may be viewed as a question about generic expansion. Given an L-theory T, we may view T as an L'-theory, for any language L' containing L. An old theorem of Winkler shows that if T is model complete and eliminates the quantifier $\exists^\infty$, then T, considered as an L'-theory, has a model companion [Win75]. We call this model companion the *generic L'-expansion* of T. When T is the theory of an infinite set in the empty language and L' is any language, the the model companion of T, considered as an L'-theory, is exactly the model companion of the empty theory in the language L'. Consequently, Jeřábek's question may be viewed as a special case of the following problem: given a model complete NSOP_1 L-theory T and language L' containing L, is the model companion of T as an L'-theory also NSOP_1?

We answer this question positively. When Winkler introduced the generic expanion of T, he also showed that if T is model-complete and eliminates the quantifier $\exists^\infty$, then there is a generic expansion of T to a theory with Skolem functions. We give a uniform analysis of Kim-independence in both constructions and show that they preserve NSOP_1. This gives a positive answer to Jeřábek's question (which was independently obtained by Jeřábek) and gives a general method for constructing new NSOP_1 theories. It also shows that NSOP_1 theories that eliminate the quantifier $\exists^\infty$ may be expanded such that every type over an arbitrary set has a global invariant extension, which greatly simplifies the analysis of Kim-independence, which *a priori* only makes sense over models.

In the final chapter, we show that PAC fields whose Galois group has NSOP_1 or NSOP_2 inverse system characterize are NSOP_1 and NSOP_2, respectively. Early work of Cherlin, van den Dries, and Macintyre [CvdDM80] show how to regard the inverse system of a profinite group as a many-sorted first order structure with a sort X_n for each n consisting of cosets of open normal subroups of index $\leq n$ in a sufficiently rich language to encode the group operation on each associated quotient and the maps between them. This approach to the model theory of profinite groups was pursued further by Chatzidakis who showed that, in this language, free profinite groups and, more generally, profinite groups with the embedding property have ω-stable theory [Cha98]. Later, Chatzidakis proved that independent amalgamation in a PAC field is, in a precise sense, completely controlled by independent amalgamation in the inverse system of the Galois group [Cha17]. From this, she was able to conclude that a PAC field is NSOP_n if the theory of the inverse system of its Galois group is NSOP_n, for $n \geq 3$. As SOP_1 and SOP_2 are defined differently from the SOP_n heirarchy for $n \geq 3$, extending the result of Chatzidakis to NSOP_1 and NSOP_2 inverse systems requires new arguments. Our proof in the NSOP_1 case relates Kim-independence in the field to Kim-independence in the Galois group. This gives the ultimate form of the results in Chapter 1 that ω-free PAC fields are NSOP_1 and in Chapter 3 that Frobenius fields are NSOP_1, since ω-free PAC fields are Frobenius fields and the absolute Galois group of a Frobenius field has an ω-stable (hence NSOP_1) inverse system. The NSOP_2 case makes use of the analysis of strongly indiscernible witnesses to SOP_2 begun in Chapter 1.

Part II

Syntax

Chapter 1

Model-theoretic tree properties

This chapter is joint work with Artem Chernikov.

1.1 Introduction

One of the central tasks of abstract model theory is to understand what kinds of complete first-order theories there are and how complicated they can be. In practice, this is achieved by classifying theories according to the combinatorial configurations that do or do not appear among the definable sets in their models. The most meaningful of these configurations, the so-called *dividing lines*, have the property that their absence signals the existence of some positive structure, while their presence indicates some kind of complexity. Dividing lines come in two flavors: local properties, which describe the combinatorics of sets defined by instances of a single formula, and global properties, which describe the interaction of definable sets generally. Stability, simplicity, NIP are examples of the former, while ω-stability, supersimplicity, and strong dependence are examples of the latter (see e.g. [Con]).

In this chapter, we study some questions around Shelah's tree property TP and its relatives SOP_1, TP_1, TP_2 and weak k-TP_1, as well as their global analogues detected by the cardinal invariants $\kappa_{\text{cdt}}(T)$, $\kappa_{\text{inp}}(T)$, and $\kappa_{\text{sct}}(T)$. Our point of departure is the third chapter of Shelah's *Classification Theory*. There, Shelah investigates the global combinatorics of stable theories in terms of a cardinal invariant $\kappa(T)$ quantifying the complexity of forking in models of T. In the final section of this chapter, he introduces variations on $\kappa(T)$ with the invariants $\kappa_{\text{cdt}}(T)$, $\kappa_{\text{sct}}(T)$, and $\kappa_{\text{inp}}(T)$ and proves several results about how they relate. In contemporary language, these invariants bound the size of approximations to the tree property, the tree property of first kind, and the tree property of the second kind consistent with T, respectively. Later as the theory developed, a property of stable theories that forking satisfies *local character* was isolated and theories satisfying this condition, the *simple theories*, were intensively studied [She80]. These theories are exactly the theories without the tree property, which is to say those theories with $\kappa_{\text{cdt}}(T)$ bounded. Nonetheless, until recently, the aforementioned invariants have received very little attention and many basic questions remain unaddressed.

Here, we focus on two such questions. Shelah proved that a theory has the tree property if and only if it has the tree property of the first kind or the tree property of the second kind [She90]. In terms of the invariants, this amounts to the assertion that $\kappa_{\text{cdt}}(T) = \infty$ if and only if $\kappa_{\text{inp}}(T) + \kappa_{\text{sct}}(T) = \infty$. It is natural to ask if this relationship persists when $\kappa_{\text{cdt}}(T)$ is bounded — in other words, if the equality $\kappa_{\text{cdt}}(T) = \kappa_{\text{inp}}(T) + \kappa_{\text{sct}}(T)$ holds in general. Shelah also proved that $\kappa_{\text{cdt}}(T) = \kappa$ is always witnessed by a sequence of formulas in a single free variable when κ is an infinite cardinal or ∞. Recently, Chernikov proved an analogous result for $\kappa_{\text{inp}}(T)$ [Che14]. We consider here whether or not the computation of $\kappa_{\text{sct}}(T)$ similarly reduces to a single free variable. These questions were both raised by Shelah (Question 7.14 in [She90]).

We do not give a complete answer to any of them, but for each of these questions there are two model-theoretically natural special cases to consider: first, the case of countable theories and, secondly, the case where one or more of the invariants in question are unbounded (which reduces to a question about configurations in a single formula). In Section 1.3, we show that $\kappa_{\text{cdt}}(T) = \kappa_{\text{inp}}(T) + \kappa_{\text{sct}}(T)$ for countable T. In Section 1.4, we show that if $\kappa_{\text{sct}}(T) = \infty$ then this will be witnessed by a formula in a single free variable by showing that TP_1 is always witnessed by a formula in one free variable. The main ingredient in our argument is the notion of a *strongly indiscernible tree*, which is more easily manipulated than the s-indiscernible trees used in other studies of the tree property of the first kind.

At the present state of the theory, the class of non-simple theories without the strict order property is poorly understood even at the level of syntax. In their study of the order $\trianglelefteq^*$, Dzamonja and Shelah introduced a weakening of TP_1 called SOP_1 [DS04]. Subsequently, Kim and Kim introduced two infinite families of properties called k-TP_1 and *weak* k-TP_1 for $k \geq 2$ and showed

$$\text{TP}_1 \iff k\text{-TP}_1 \iff \text{weak 2-TP}_1 \implies \text{weak 3-TP}_1 \implies \ldots \implies \text{SOP}_1$$

It was left open whether the properties weak k-TP_1 are inequivalent for distinct k and whether or not weak k-TP_1 is equivalent to TP_1 [KK11]. In our work on proving that TP_1 is witnessed by a formula in one free variable, we obtained unexpectedly a simple and direct proof that the weak k-TP_1 hierarchy collapses and that they are all equivalent to TP_1.

In the final two sections of the chapter, we study theories without the property SOP_1. We show that independent amalgamation fails in a strong way in theories with SOP_1 and that they are in fact characterized by this feature. This gives rise to a useful criterion for showing that a theory is NSOP_1 (and hence NTP_1). Leveraging work of Granger [Gra99] and Chatzidakis [Cha02], this allows us to conclude that both the two sorted theory of infinite-dimensional vector spaces over an algebraically closed field with a generic bilinear form, as well as the theory of ω-free PAC fields of characteristic zero are NSOP_1. Finally, we generalize the construction of the theory of parametrized equivalence relations T^*_{feq} to give a general method for constructing NSOP_1 theories from simple ones. We learned after this work was completed that essentially the same construction had been studied by Baudisch [Bau02], but our emphasis is different. We show that the independence theorem holds for these structures, allowing us to obtain a proof that T^*_{feq} is NSOP$_1$ as a corollary.

1.2 Preliminaries on indiscernible trees

We fix a complete first-order theory T in a language L, $\mathbb{M} \models T$ is a monster model. In several of the arguments below, we will make use of the notion of an indiscernible tree. For our purposes, there are two different languages we will need to place on the index model: $L_{s,\lambda} = \{\triangleleft, \wedge, <_{lex}, (P_\alpha : \alpha < \lambda)\}$ and $L_0 = \{\triangleleft, \wedge, <_{lex}\}$ where λ is a cardinal. We may view the tree $\kappa^{<\lambda}$ as an $L_{s,\lambda}-$ or L_0-structure in a natural way, interpreting $\triangleleft$ as the tree partial order, $\wedge$ as the binary meet function, $<_{lex}$ as the lexicographic order, and P_α as a predicate which identifies the αth level (we will only consider $\kappa = 2$ and $\kappa = \omega$). See [KKS14] and [TT12] for more details.

Definition 1.2.1. Suppose that $(a_\eta)_{\eta \in \kappa^{<\lambda}}$ and $(a_{\alpha,i})_{\alpha<\kappa, i<\omega}$ are collections of tuples and C is a set of parameters in some model.

1. We say $(a_\eta)_{\eta \in \kappa^{<\lambda}}$ is an *s-indiscernible tree over* C if
$$\text{qftp}_{L_{s,\lambda}}(\eta_0, \ldots, \eta_{n-1}) = \text{qftp}_{L_{s,\lambda}}(\nu_0, \ldots, \nu_{n-1})$$
implies $\text{tp}(a_{\eta_0}, \ldots, a_{\eta_{n-1}}/C) = \text{tp}(a_{\nu_0}, \ldots, a_{\nu_{n-1}}/C)$, for all $n \in \omega$.

2. We say $(a_\eta)_{\eta \in \kappa^{<\lambda}}$ is a *strongly indiscernible tree over* C if
$$\text{qftp}_{L_0}(\eta_0, \ldots, \eta_{n-1}) = \text{qftp}_{L_0}(\nu_0, \ldots, \nu_{n-1})$$
implies $\text{tp}(a_{\eta_0}, \ldots, a_{\eta_{n-1}}/C) = \text{tp}(a_{\nu_0}, \ldots, a_{\nu_{n-1}}/C)$, for all $n \in \omega$.

3. We say $(a_{\alpha,i})_{\alpha<\kappa, i<\lambda}$ is a *mutually indiscernible array over* C if, for all $\alpha < \kappa$, $(a_{\alpha,i})_{i<\lambda}$ is a sequence indiscernible over $C \cup \{a_{\beta,j} : \beta < \kappa, \beta \neq \alpha, j < \lambda\}$.

Lemma 1.2.2. *Let $(a_\eta : \eta \in \kappa^{<\lambda})$ be a tree strongly indiscernible over a set of parameters C.*

1. *All paths have the same type over C: for any $\eta, \nu \in \kappa^\lambda$, $\text{tp}((a_{\eta|\alpha} : \alpha < \lambda)/C) = \text{tp}((a_{\nu|\alpha} : \alpha < \lambda)/C)$.*

2. *For any $\eta \perp \nu \in \kappa^{<\lambda}$ and any ξ, $\text{tp}(a_\eta, a_\nu/C) = \text{tp}(a_{\xi\frown 0}, a_{\xi\frown 1}/C)$.*

3. *The tree $(a_{0\frown\eta} : \eta \in \kappa^{<\lambda})$ is strongly indiscernible over $a_\emptyset C$.*

Proof. (1) This follows by strong indiscernibility of the tree as for any $\eta, \nu \in \kappa^{<\lambda}$, $\text{qftp}_{L_0}((\eta|\alpha : \alpha < \lambda)) = \text{qftp}_{L_0}((\nu|\alpha : \alpha < \lambda))$.

(2) Let $\eta \perp \nu \in \kappa^{<\lambda}$ be given, without loss of generality $\eta <_{lex} \nu$ and let $\mu = \eta \wedge \nu$. Then there are $i < j < \kappa$ so that $\mu \frown \langle i \rangle \trianglelefteq \eta$ and $\mu \frown \langle j \rangle \trianglelefteq \nu$. Then $\mathrm{qftp}_{L_0}(\eta, \nu) = \mathrm{qftp}_{L_0}(\mu \frown \langle i \rangle, \mu \frown \langle j \rangle) = \mathrm{qftp}_{L_0}(\mu \frown 0, \mu \frown 1) = \mathrm{qftp}_{L_0}(\xi \frown 0, \xi \frown 1)$, and we conclude by strong indiscernibility of the tree.

(3) Clear as $\mathrm{qftp}_{L_0}(\bar{\eta}) = \mathrm{qftp}_{L_0}(\bar{\nu})$ implies $\mathrm{qftp}_{L_0}(\bar{\eta}, \emptyset) = \mathrm{qftp}_{L_0}(\bar{\nu}, \emptyset)$, provided $\emptyset$ is not enumerated in neither $\bar{\eta}$ nor $\bar{\nu}$. $\square$

Lemma 1.2.3. *Let $(a_\eta : \eta \in \kappa^{<\lambda})$ be a tree s-indiscernible over a set of parameters C.*

1. *All paths have the same type over C: for any $\alpha, \nu \in \kappa^\lambda$, $\mathrm{tp}((a_{\eta|\alpha})_{\alpha<\lambda}/C) = \mathrm{tp}((a_{\nu|\alpha})_{\alpha<\lambda}/C)$.*
2. *Suppose $\{\eta_\alpha : \alpha < \gamma\} \subseteq \kappa^{<\lambda}$ satisfies $\eta_\alpha \perp \eta_{\alpha'}$ whenever $\alpha \neq \alpha'$. Then the array $(b_{\alpha,\beta})_{\alpha<\gamma,\beta<\kappa}$ defined by*
$$b_{\alpha,\beta} = a_{\eta_\alpha \frown \langle \beta \rangle}$$
is mutually indiscernible over C.

Proof. (1) This follows by s-indiscernibility of the tree as for any $\eta, \nu \in \kappa^{<\lambda}$, $\mathrm{qftp}_{L_s}((\eta|\alpha : \alpha < \lambda)) = \mathrm{qftp}_{L_s}((\nu|\alpha : \alpha < \lambda))$.

(2) Fix $\alpha < \gamma$ and let $A = \{a_{\eta_{\alpha'} \frown \langle \beta \rangle} : \alpha \neq \alpha' < \gamma, \beta < \kappa\} \cup C$. As the elements of $\{\eta_\alpha : \alpha < \gamma\}$ are pairwise incomparable, it is easy to check that for any $\beta_0 < \ldots < \beta_{n-1} < \kappa$ and $\beta'_0 < \ldots < \beta'_{n-1} < \kappa$,

$$\mathrm{qftp}_{L_s}(a_{\eta_\alpha \frown \langle \beta_0 \rangle}, \ldots, a_{\eta_\alpha \frown \langle \beta_{n-1} \rangle}/A) = \mathrm{qftp}_{L_s}(a_{\eta_\alpha \frown \langle \beta'_0 \rangle}, \ldots, a_{\eta_\alpha \frown \langle \beta'_{n-1} \rangle}/A),$$

which proves (2). $\square$

Now we note that s-indiscernible and strongly indiscernible trees exist.

Definition 1.2.4. Suppose I is an L'-structure, where L' is some language. We say that I-indexed indiscernibles have the *modeling property* if, given any $(a_i : i \in I)$ from $\mathbb{M}$, there is an I-indexed indiscernible $(b_i : i \in I)$ in $\mathbb{M}$ *locally based* on the (a_i): given any finite set of formulas Δ from L and a finite tuple $(t_0, \ldots, t_{n-1})$ from I, there is a tuple $(s_0, \ldots, s_{n-1})$ from I so that
$$\mathrm{qftp}_{L'}(t_0, \ldots, t_{n-1}) = \mathrm{qftp}_{L'}(s_0, \ldots, s_{n-1})$$
and also
$$\mathrm{tp}_\Delta(b_{t_0}, \ldots, b_{t_{n-1}}) = \mathrm{tp}_\Delta(a_{s_0}, \ldots, a_{s_{n-1}}).$$

Fact 1.2.5. [TT12, Sco12, KKS14] Let I_0 denote the L_0-structure $(\omega^{<\omega}, \trianglelefteq, <_{lex}, \wedge)$ and I_s be the $L_{s,\omega}$-structure $(\omega^{<\omega}, \trianglelefteq, <_{lex}, \wedge, (P_\alpha)_{\alpha<\omega})$ with all symbols being given their intended interpretations and each P_α naming the elements of the tree at level α. Then strongly indiscernible trees (I_0-indexed indiscernibles) and s-indiscernible trees (I_s-indexed indiscernibles) have the *modeling property.*

In the arguments below, we will often argue by induction where at each stage it is necessary to modify a tree of tuples in a way that maintains the indiscernibility of the tree. A convenient way of organizing these arguments is to make a catalogue of operations on indiscernible trees and prove that these operations preserve the relevant indiscernibility.

Definition 1.2.6. Fix $k \geq 1$.

1. (widening) The *k-fold widening of $(a_\eta)_{\eta\in\omega^{<\omega}}$ at level n* is defined to be the tree $(a'_\eta)_{\eta\in\omega^{<\omega}}$ where

$$a'_\eta = \begin{cases} a_\eta & \text{if } l(\eta) < n \\ (a_{\nu\frown(ki)\frown\xi}, \ldots, a_{\nu\frown(ki+(k-1))\frown\xi}) & \text{if } \eta = \nu \frown i \frown \xi \\ & \text{where } \nu \in \omega^{n-1}, i \in \omega, \xi \in \omega^{<\omega}. \end{cases}$$

2. (stretching) The *k-fold stretch of $(a_\eta)_{\eta\in\omega^{<\omega}}$ at level n* is defined to be the tree $(a''_\eta)_{\eta\in\omega^{<\omega}}$ where

$$a''_\eta = \begin{cases} a_\eta & \text{if } l(\eta) < n \\ (a_\eta, a_{\eta\frown 0}, \ldots, a_{\eta\frown 0^{k-1}}) & \text{if } l(\eta) = n \\ a_{\nu\frown 0^{k-1}\frown\xi} & \text{if } \eta = \nu \frown \xi \text{ for } \nu \in \omega^n, \xi \neq \emptyset \end{cases}$$

3. (fattening) Given a tree $(a_\eta)_{\eta\in 2^{<\kappa}}$, define the *$k$-fold fattening of $(a_\eta)_{\eta\in 2^{<\kappa}}$* to be the tree $(a^{(k)}_\eta)_{\eta\in 2^{<\kappa}}$ by induction as follows: for each $\eta \in 2^{<\kappa}$ let $a^{(0)}_\eta = a_\eta$. If $(a^{(n)}_\eta)_{\eta\in 2^{<\kappa}}$ has been defined, for each $\eta \in 2^{<\kappa}$, let $a^{(n+1)}_\eta = (a^{(n)}_{0\frown\eta}, a^{(n)}_{1\frown\eta})$. Let $C_k = \{a_\eta : \eta \in 2^{<k}\}$, the *stump below k*. Set $C_0 = \emptyset$.

4. (restricting) Given the tree $(a_\eta)_{\eta\in\lambda^{<\kappa}}$ and $W \subseteq \kappa$, we define the *restriction of $(a_\eta)_{\eta\in\lambda^{<\kappa}}$ to W* to be the collection of tuples

$$\{a_\eta : l(\eta) \in W \text{ and if } \beta \notin W, \text{ then } \eta(\beta) = 0\}.$$

   If the order type of W is α, the restriction of $(a_\eta)_{\eta\in\lambda^{<\kappa}}$ may be naturally identified with $(a_\eta)_{\eta\in\lambda^{<\alpha}}$.

5. (elongating) Given $\eta \in \kappa^{<\omega}$, with $l(\eta) = n$, define $\tilde{\eta} \in \kappa^{<\omega}$ to be the tuple with length $k(l(\eta)-1)+1$ defined by

$$\tilde{\eta}(i) = \begin{cases} \eta(i/k) & \text{if } k | i \\ 0 & \text{otherwise} \end{cases}$$

   Then define the *k-fold elongation* of $(a_\eta)_{\eta\in\kappa^{<\omega}}$ to be the tree $(b_\eta)_{\eta\in\kappa^{<\omega}}$ where

$$b_\eta = (a_{\tilde{\eta}}, a_{\tilde{\eta}\frown 0}, \ldots, a_{\tilde{\eta}\frown 0^{k-1}}).$$

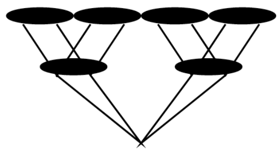

2-fold widening at level 1

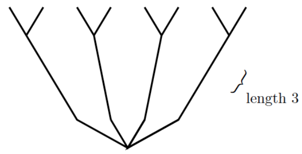

3-fold stretch at level 1

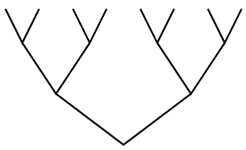

0-fold fattening

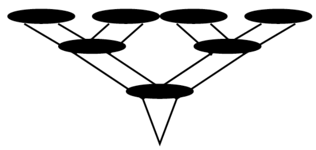

1-fold fattening

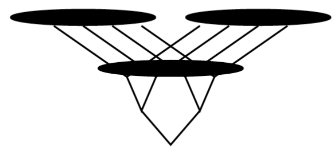

2-fold fattening

Proposition 1.2.7. *1. s-indiscernibility is preserved under widening, stretching, fattening, restriction, and elongating.*

2. Strong indiscernibility is preserved under restriction, fattening, and elongating. Moreover, if $(a_\eta)_{\eta\in 2^{<\omega}}$ is strongly indiscernible, then the k-fold fattening $(a^{(k)})_{\eta\in 2^{<\omega}}$ is strongly indiscernible over C_k.

Proof. The proofs of these facts can be found in Section 1.7. $\square$

1.3 Cardinal invariants and tree properties

Definition 1.3.1. Suppose T is a complete theory and $\varphi(x;y) \in L$ is a formula in the language of T.

1. $\varphi(x;y)$ has the *tree property* (TP) if there is $k<\omega$ and a tree of tuples $(a_\eta)_{\eta\in\omega^{<\omega}}$ in $\mathbb{M}$ such that
   - for all $\eta\in\omega^\omega$, $\{\varphi(x;a_{\eta|\alpha}) : \alpha<\omega\}$ is consistent,
   - for all $\eta\in\omega^{<\omega}$, $\{\varphi(x;a_{\eta\frown\langle i\rangle}) : i<\omega\}$ is k-inconsistent.
2. $\varphi(x;y)$ has the *tree property of the first kind* (TP_1) if there is a tree of tuples $(a_\eta)_{\eta\in\omega^{<\omega}}$ in $\mathbb{M}$ such that
   - for all $\eta\in\omega^\omega$, $\{\varphi(x;a_{\eta|\alpha}) : \alpha<\omega\}$ is consistent,
   - for all $\eta\perp\nu$ in $\omega^{<\omega}$, $\{\varphi(x;a_\eta),\varphi(x;a_\nu)\}$ is inconsistent.
3. $\varphi(x;y)$ has the *tree property of the second kind* (TP_2) if there is a $k<\omega$ and an array $(a_{\alpha,i})_{\alpha<\omega,i<\omega}$ in $\mathbb{M}$ such that
   - for all functions $f:\omega\to\omega$, $\{\varphi(x;a_{\alpha,f(\alpha)}) : \alpha<\omega\}$ is consistent,

- for all α, $\{\varphi(x; a_{\alpha,i}) : i < \omega\}$ is k-inconsistent.

4. T has one of the above properties if some formula does modulo T.

It is easy to see that if a theory has the tree property of the first or second kind, then it also has the tree property. Remarkably, the converse is also true.

Fact 1.3.2. [She90] A complete theory T has TP if and only if it has TP_1 or TP_2.

The above theorem was first proven in different language, before any of the three properties were actually defined. The purpose of this section is to prove a refinement of this theorem, by studying the relationship between approximations to the tree property and those to the tree property of the first or second kind. In order to do so, however, it will be necessary to return to the vocabulary in which Fact 1.3.2 was initially formulated.

Definition 1.3.3. The following notions were introduced in [She90].

1. A cdt-*pattern of depth* κ is a sequence of formulas $\varphi_i(x; y_i)$ $(i < \kappa, i$ successor) and numbers $n_i < \omega$, and a tree of tuples $(a_\eta)_{\eta \in \omega^{<\kappa}}$ for which

   a) $p_\eta = \{\varphi_i(x; a_{\eta|i}) : i \text{ successor }, i < \kappa\}$ is consistent for $\eta \in \omega^\kappa$,

   b) $\{\varphi_i(x; a_{\eta^\frown\langle\alpha\rangle}) : \alpha < \omega, i = l(\eta) + 1\}$ is n_i-inconsistent.

   A cdt-pattern with $n_i \leq n$ for all $i < \kappa$, is called a (cdt, n)-pattern.

2. An inp-*pattern of depth* κ is a sequence of formulas $\varphi_i(x; y_i)$ $(i < \kappa)$, sequences $(a_{i,\alpha} : \alpha < \omega)$, and numbers $n_i < \omega$ such that

   a) for any $\eta \in \omega^\kappa$, $\{\varphi_i(x; a_{i,\eta(i)}) : i < \kappa\}$ is consistent,

   b) for any $i < \kappa$, $\{\varphi_i(x; a_{i,\alpha}) : \alpha < \omega\}$ is n_i-inconsistent.

3. An sct-*pattern of depth* κ is a sequence of formulas $\varphi_i(x; y_i)$ $(i < \kappa)$ and a tree of tuples $(a_\eta)_{\eta \in \omega^{<\kappa}}$ such that

   a) for every $\eta \in \omega^\kappa$, $\{\varphi_\alpha(x; a_{\eta|\alpha}) : 0 < \alpha < \kappa, \alpha \text{ successor}\}$ is consistent,

   b) If $\eta \in \omega^\alpha$, $\nu \in \omega^\beta$, α, β are successors, and $\nu \perp \eta$ then the formulas $\{\varphi_\alpha(x; a_\eta), \varphi_\beta(x; a_\nu)\}$ are inconsistent.

   If instead of (b), we have: for any pairwise incomparable $(\eta_i : i < k)$, $\{\varphi_{l(\eta_i)}(x; a_{\eta_i}) : i < k\}$ is inconsistent, then we call this a (sct, k)-pattern.

4. For $X \in \{\mathrm{cdt}, \mathrm{sct}, \mathrm{inp}\}$, we define $\kappa^n_X(T)$ to be the first cardinal κ so that there is no X-pattern of depth κ in n free variables, and ∞ if no such κ exists. We define $\kappa_X(T) = \sup_{n\in\omega}\{\kappa^n_X\}$.

Remark 1.3.4. We note that the notion of a (cdt, n)-pattern *strengthens* that of a cdt-pattern by imposing a uniform finite bound on the size of the inconsistency at each level, while the notion of an (sct, n)-pattern *weakens* that of an sct-pattern by only requiring any n incomparable elements to be inconsistent rather than any 2. One can regard an (sct, n)-pattern as an approximation to a witness to n-TP_1 (see Definition 1.4.1 below).

Observation 1.3.5. Fix a complete theory T.

1. $\kappa^n_{\mathrm{sct}}(T) \geq n$, $\kappa^n_{\mathrm{inp}}(T) \geq n$ and $\kappa^n_{\mathrm{cdt}}(T) \geq n$ for all n.
2. a) $\kappa_{\mathrm{cdt}}(T) = \infty$ if and only if $\kappa_{\mathrm{cdt}}(T) > |T|^+$ if and only if T has TP.
   b) $\kappa_{\mathrm{sct}}(T) = \infty$ if and only if $\kappa_{\mathrm{sct}}(T) > |T|^+$ and only if T has TP_1.
   c) $\kappa_{\mathrm{inp}}(T) = \infty$ if and only if $\kappa_{\mathrm{inp}}(T) > |T|^+$ if and only if T has TP_2.
3. $\max\{\kappa^n_{\mathrm{sct}}(T), \kappa^n_{\mathrm{inp}}(T)\} \leq \kappa^n_{\mathrm{cdt}}(T)$.

Proof. (1) follows from the fact that "=" is in the language.

(2) As each case is entirely similar, we'll sketch the argument for (a) only. If $\kappa_{\mathrm{cdt}}(T) > |T|^+$, then in the pattern witnessing it we may assume that $\varphi_i(x, y_i) = \varphi(x, y)$ and $k_i = k$, because $|T| \geq \aleph_0$. This is a witness to TP. And then using compactness we can find a pattern witnessing that $\kappa^n_{\mathrm{cdt}}(T) > \kappa$ for any cardinal κ.

(3) If $\varphi_i(x; y_i)$ $(i < \kappa)$, $(a_{i,\alpha} : \alpha < \omega)$, $(n_i)_{i<\omega}$ form an inp-pattern of depth κ, obtain a cdt-pattern of depth κ with respect to the same formulas by defining $(b_\eta)_{\eta \in \omega^{<\kappa}}$ by $b_\eta = a_{l(\eta), \eta(l(\eta)-1)}$. □

Lemma 1.3.6. *(1) If there is an sct-pattern (cdt-pattern) of depth κ modulo T, then there is an sct-pattern (cdt-pattern) $\varphi_\alpha(x; y_\alpha)$, $(a_\eta)_{\eta \in \omega^{<\kappa}}$ in the same number of free variables so that $(a_\eta)_{\eta \in \omega^{<\kappa}}$ is an s-indiscernible tree.*

(2) If there is an inp-pattern of depth κ modulo T, then there is an inp-pattern $\varphi_\alpha(x; y_\alpha)$ $(\alpha < \kappa)$, $(k_\alpha)_{\alpha<\kappa}$, $(a_{\alpha,i})_{\alpha<\kappa, i<\omega}$ in the same number of free variables so that $(a_{\alpha,i})_{\alpha<\kappa, i<\omega}$ is a mutually indiscernible array.

Proof. (1) By compactness and Fact 3.5.4.

(2) This is Lemma 2.2 of [Che14]. □

Now we fix a complete theory T and for $X \in \{\mathrm{cdt}, \mathrm{sct}, \mathrm{inp}\}$, we write κ_X for $\kappa_X(T)$.

Proposition 1.3.7. *Assume that $\kappa^n_{\mathrm{cdt}} \geq \aleph_0$. Then either $\kappa^n_{\mathrm{inp}} \geq \aleph_0$ or $\kappa^n_{\mathrm{sct},k} \geq \aleph_0$ for some $k \in \omega$ (i.e. there are $(\kappa_{\mathrm{sct}}, k)$-patterns in n variables of arbitrary finite depth). In fact, if $\kappa^n_{\mathrm{inp}} < \aleph_0$, then one can take $k = \kappa^n_{\mathrm{inp}}$.*

Proof. If $\kappa^n_{\mathrm{inp}} \geq \aleph_0$ does not hold, then in fact we have $\kappa^n_{\mathrm{inp}} \leq k$ for some $k \in \omega$.

Fix an arbitrary $m \in \omega$, then by assumption and Lemma 1.3.6 we can find $(a_\eta : \eta \in \omega^{<2m}), (\varphi_i(x, y_i) : i <$
an s-indiscernible cdt-pattern with $|x| = n$, i.e.:

1. $(a_\eta : \eta \in \omega^{<2m})$ is an s-indiscernible tree,
2. $\{\varphi_i(x, a_{\eta\restriction i}) : i < 2m\}$ is consistent for every $\eta \in \omega^{2m}$,
3. $\left\{\varphi_i\left(x, a_{\eta\frown\langle j\rangle}\right) : j \in \omega\right\}$ is k_i-inconsistent for every $i < 2m-1$ and $\eta \in \omega^i$.

For $l < m$ and $\nu \in \omega^l$ we define $\nu^* = (\nu(0), 0, \nu(1), 0, \ldots, \nu(l-1), 0) \in \omega^{<2m}$. Let $\{\nu_0, \ldots, \nu_{k-1}\} \subseteq \omega^{<m}$ be pairwise $\trianglelefteq$-incomparable, and let $l_i = l(\nu_i^*)$.

Claim. $\left\{\varphi_{l_i}\left(x, a_{\nu_i^*}\right) : i < k\right\}$ is inconsistent.

Proof. By definition of ν_i^* and assumption on ν_i's it follows that for any $i, i' < k$ the elements $\nu_i^* \restriction (l_i - 1)$ and $\nu_{i'}^* \restriction (l_{i'} - 1)$ are incomparable. Then by Lemma 2.2.4(2) we see that the sequences $\bar{a}_i = \left(a_{\nu_i^*\restriction(l_i-1)\frown\langle j\rangle} : j \in \omega\right)$ are mutually indiscernible. But if $\left\{\varphi_{l_i}\left(x, a_{\nu_i^*}\right) : i < k\right\}$ was consistent, this would give us an inp-pattern of depth k, contrary to the assumption (as $\left\{\varphi_{l_i}\left(x, a_{\nu_i^*\restriction(l_i-1)\frown\langle j\rangle}\right) : j \in \omega\right\}$ is k_{l_i}-inconsistent for every i).

Now using the claim it is easy to see that $\left\{\varphi_{2l(\eta)}(x, a_{\eta^*}) : \eta \in \omega^{<m}\right\}$ is an (sct, k)-pattern of depth m. As m was arbitrary, we conclude that $\kappa^n_{\mathrm{sct},k} \geq \aleph_0$. ☐

Proposition 1.3.8. *Let $k < \omega$ be fixed. Assume that for any $n < \omega$ we have, in some fixed number of variables, an (sct, k)-pattern of depth n. Then there are, in the same number of variables, $(\mathrm{cdt}, 2)$-patterns of arbitrary finite depth.*

Proof. Let $m \in \omega$ be arbitrary, and let $(a_\eta : \eta \in \omega^{<m\times m}), (\varphi_i(x, y_i) : i < m \times m)$ be an s-indiscernible (sct, k)-pattern - in particular this is a cdt-pattern such that for $i < m \times m$, $\{\varphi_i(x; a_\eta) : l(\eta) = i\}$ is k-inconsistent.

For $i < m$, consider

$$\Gamma_i(x) = \bigwedge_{l<m} \left(\varphi_{i\times m+l}\left(x, a_{0^{i\times m}\frown 0\frown 0^{l-1}}\right) \wedge \varphi_{i\times m+l}\left(x, a_{0^{i\times m}\frown 1\frown 0^{l-1}}\right)\right).$$

Case 1. $\Gamma_i(x)$ is consistent for some $i < m$.

Obtain an s-indiscernible tree, using Lemma 1.2.7(1), by first taking the 2-fold widening of $(a_\eta)_{\eta\in\omega^{m\times m}}$ at level $i \times m + 1$, then taking the restriction to $\{i \times m + l : l < m\}$. Let $(\psi_l : l < m)$ be chosen so that

$$\psi_l(x, b_{0^l}) = \varphi_{i\times m+l}\left(x, a_{0^{i\times m}\frown 0\frown 0^{l-1}}\right) \wedge \varphi_{i\times m+l}\left(x, a_{0^{i\times m}\frown 1\frown 0^{l-1}}\right).$$

Then $(b_\eta : \eta \in \omega^{<m}), (\psi_l : l < m)$ is a cdt-pattern of depth m such that, for all $l < m$, $\{\psi_l(x; b_\eta) : l(\eta) = l\}$ is $\lfloor \frac{k}{2} \rfloor$-inconsistent.

Case 2. $\Gamma_i(x)$ is inconsistent for every $i < m$.

Using Lemma 1.2.7(1), obtain an s-indiscernible tree $(b_\eta)_{\eta\in\omega^{<m}}$ by taking the m-fold elongation of $(a_\eta)_{\eta\in\omega^{<m\times m}}$. Let $(\psi_l : l < m)$ be chosen so that

$$\psi_l(x; b_{0^l}) = \bigwedge_{r<m} \varphi_{l\times m+r}(x; a_{0^{l\times m}\frown 0^r}).$$

Then $(b_\eta)_{\eta\in\omega^{<m}}$, $(\psi_l : l < m)$ is an $(\mathrm{cdt}, 2)$-pattern.

Repeating several times if necessary we conclude. ☐

For $\kappa \leq \omega$, finding an sct-pattern of depth κ is equivalent to finding a $(\mathrm{cdt}, 2)$-pattern of depth κ.

Lemma 1.3.9. *Let $\kappa \leq \omega$, and let $(a_\eta : \eta \in \omega^{<\kappa}), (\varphi_i(x, y_i) : i < \kappa)$ be a $(\mathrm{cdt}, 2)$-pattern (i.e. for every $\eta \in \omega^{<\kappa}$ the set $\left\{\varphi_{l(\eta)+1}(x, a_{\eta\frown j}) : j \in \omega\right\}$ is 2-inconsistent). For $\eta \in \omega^{<\kappa}$ define $b_\eta = a_{\eta\restriction 0} a_{\eta\restriction 1} \dots a_{\eta\restriction(l(\eta)-1)} a_\eta$ and $\psi_i(x; y_{i,0}, \dots, y_{i,i-1}) = \bigwedge_{j<i} \varphi_j(x, y_j)$. Then $(b_\eta : \eta \in \omega^{<\kappa}), (\psi_i(x, \bar{y}_i) : i < \kappa)$ is an* sct*-pattern.*

Proof. If $\eta \in \omega^n$ for $n < \kappa$, then the set $\{\psi_i(x, b_{\eta\restriction i}) : i < n\}$ contains only conjunctions of formulas from $\{\varphi_i(x, a_{\eta\restriction i}) : i < n\}$ which is consistent by assumption. On the other hand if $\eta_1, \eta_2 \in \omega^{<\kappa}$ are incomparable, let $\eta = \eta_1 \wedge \eta_2$. Then $\psi_{l(\eta_1)}(x, b_{\eta_1})$ implies $\varphi_{l(\eta)+1}\left(x, a_{\eta\frown\eta_1(l(\eta)+1)}\right)$ and $\psi_{l(\eta_2)}(x, b_{\eta_2})$ implies $\varphi_{l(\eta)+1}\left(x, a_{\eta\frown\eta_2(l(\eta)+1)}\right)$, and these two implied formulas are inconsistent by assumption. $\square$

Combining Propositions 1.3.7 and 1.3.8 with Lemma 1.3.9, we have:

Proposition 1.3.10. *If $\kappa^n_{\mathrm{cdt}} \geq \aleph_0$, then either $\kappa^n_{\mathrm{inp}} \geq \aleph_0$ or $\kappa^n_{\mathrm{sct}} \geq \aleph_0$.*

Remark 1.3.11. Inspecting the proof, we actually get the following bound: $\kappa^n_{\mathrm{sct}} \geq \left(\frac{\kappa^n_{\mathrm{cdt}}}{2}\right)^{\frac{1}{\kappa^n_{\mathrm{inp}}}}$.

The next proposition is an analog of Proposition 1.3.8 for inp-patterns. It is not used in this chapter, but we include it for reference.

Proposition 1.3.12. *Let $k < \omega$ be fixed. Assume that for any $n < \omega$ we have, in some fixed number of free variables, an* inp*-pattern of depth n such that each row is k-inconsistent. Then there are, in the same number of variables,* inp*-patterns of arbitrary finite depths in which every row is 2-inconsistent.*

Proof. Let $m \in \omega$ be arbitrary, and let $(a_{i,j})_{i<m\times m, j\in\omega}, (\varphi_i(x, y_i))_{i<m\times m}$ be an inp-pattern with mutually indiscernible rows such that every row is k-inconsistent. For $i < m$, consider $\Gamma_i(x) = \bigwedge_{i\times m \leq l < (i+1)\times m} (\varphi_l(x, a_{l,0}) \wedge \varphi_l(x, a_{l,1}))$.

Case 1. $\Gamma_i(x)$ is consistent for some $i < m$.

Then for $l < m$ we take $\psi_l(x, b_{l,0}) = \varphi_{i\times m+l}(x, a_{i\times m+l,0}) \wedge \varphi_{i\times m+l}(x, a_{i\times m+l,1})$ and $b_{l,j} = a_{i\times m+l,2j} a_{i\times m+l,2j+1}$.

Case 2. $\Gamma_i(x)$ is inconsistent for every $i < m$.

Then for $l < m$ we take $\psi_l(x, b_{l,0}) = \bigwedge_{r<m} \varphi_{l\times m+r}(x, a_{l\times m+r,0})$ and $b_{l,j} = (a_{l\times m+r,j} : r < m)$.

It is easy to see that in each of the cases $(b_{i,j})_{i<m,j<\omega}, (\psi_i(x, y_i))_{i<m}$ is an inp-pattern of depth m, and moreover it is $\max\left\{2, \left\lceil\frac{k}{2}\right\rceil\right\}$-inconsistent ($\left\lceil\frac{k}{2}\right\rceil$-inconsistent in the first case and 2-inconsistent in the second case). As m was arbitrary, this shows that there are inp-pattern of arbitrarily large finite depth with $\max\left\{2, \left\lceil\frac{k}{2}\right\rceil\right\}$-inconsistent rows. Repeating the argument several times if necessary we conclude. $\square$

Now we consider the case of countably infinite patterns.

Proposition 1.3.13. *$\kappa^n_{cdt} \geq \aleph_1$ implies $\kappa^n_{sct} \geq \aleph_1$.*

Proof. Suppose $(\varphi_i : i < \omega)$, $(a_\eta)_{\eta\in\omega^{<\omega}}$ is a cdt-pattern. By replacing a_η with $b_\eta = (a_\emptyset, a_{\eta|1}, \ldots, a_{\eta|l(\eta)-1}, a_\eta)$ and $\varphi_i(x; a_\eta)$ by

$$\psi_i(x; b_\eta) := \bigwedge_{j\leq i} \varphi_j(x; a_{\eta|j}),$$

if necessary, we may assume that if $\nu \lhd \eta$, then

$$\models (\forall x)[\varphi_{l(\eta)}(x; a_\eta) \to \varphi_{l(\nu)}(x; a_\nu)].$$

Then by replacing $(a_\eta)_{\eta\in\omega^{<\omega}}$ by an s-indiscernible tree locally based on it, we may moreover assume the $(a_\eta)_{\eta\in\omega^{<\omega}}$ are s-indiscernible by Fact 3.5.4.

By induction, we will construct cdt-patterns $(\varphi_i^n : i < \omega)$, $(a_\eta^n)_{\eta\in\omega^{<\omega}}$ so that

1. $(a_\eta^n)_{\eta\in\omega^{<\omega}}$ is s-indiscernible.
2. For all $\eta \in \omega^{<n}$ and $i < j$,
$$\{\varphi^n_{l(\eta)+1}(x; a^n_{\eta\frown\langle i\rangle}), \varphi^n_{l(\eta)+1}(x; a^n_{\eta\frown\langle j\rangle})\}$$
is inconsistent.
3. If $\nu \lhd \eta$, then
$$\models (\forall x)[\varphi^n_{l(\eta)}(x; a^n_\eta) \to \varphi^n_{l(\nu)}(x; a^n_\nu)].$$
4. For all η, if $n, n' \geq l(\eta)$, then $a^n_\eta = a^{n'}_\eta$. For all $m \leq m'$, $\varphi^{m'}_m = \varphi^m_m$.

For the base case, let $\varphi^0_i = \varphi_i$ for all i and $a^0_\eta = a_\eta$ for all η. (1) is satisfied by assumption, (2) is vacuous, and (3) follows from the initial remarks above. Now suppose we have constructed $(\varphi^n_i : i < \omega)$ and $(a^n_\eta)_{\eta\in\omega^{<\omega}}$. By definition of a cdt-pattern, there is a least $k \geq 1$ so that

$$\bigcup_{i<2^k} \{\varphi^n_{n+1+j}(x; a^n_{0^n\frown\langle i\rangle\frown 0^j}) : j < \omega\}$$

is inconsistent. By compactness, there is N so that

$$\bigcup_{i<2^k} \{\varphi^n_{n+1+j}(x; a_{0^n\frown\langle i\rangle\frown 0^j}) : j < N\} \tag{1.1}$$

is inconsistent. Let $(b_\eta)_{\eta\in\omega^{<\omega}}$ be the N-fold stretch of $(a^n)_{\eta\in\omega^{<\omega}}$ at level n. Let $(\psi_i(x; z_i) : i < \omega)$ be defined as follows: for $i \leq n$, $z_i = y_i$ and $\psi_i(x; z_i) = \varphi_i(x; y_i)$. Let $z_{n+1} = (y_{n+1}, y_{n+2}, \ldots, y_{n+N})$ and

$$\psi_{n+1}(x; z_{n+1}) = \bigwedge_{j<N} \varphi^n_{n+1+j}(y; y_{n+1+j}).$$

Finally, for $i > n+1$, let $z_i = y_{i+N-1}$ and $\psi_i(x; z_i) = \varphi_{i+N-1}(x; y_{i+N-1})$. By Lemma 1.7.4, $(b_\eta)_{\eta\in\omega^{<\omega}}$ is an s-indiscernible tree and, by construction, $(\psi_i(x; z_i) : i < \omega)$, $(b_\eta)_{\eta\in\omega^{<\omega}}$ is a cdt-pattern. Moreover, this cdt-pattern satisfies

5. $\{\psi_{n+1}(x; b_{0^n \frown \langle i \rangle}) : i < 2^k\}$ is inconsistent and

6. $\{\psi_{n+1+j}(x; b_{0^n \frown \langle i \rangle \frown 0^j}) : i < 2^{k-1}, j < \omega\} \cup \{\psi_l(x; b_{0^l}) : l < \omega\}$ is consistent.

Condition (5) follows by the inconsistency (1.1) and the definition of ψ_{n+1}. To see (6), we note that by the minimality of k,

$$\{\psi_{n+1+j}(x; b_{0^n \frown \langle i \rangle \frown 0^j}) : i < 2^{k-1}, j < \omega\}$$

is consistent. By (3) above and the definition of the ψ_m, this establishes (6).

Let $(c_\eta)_{\eta \in \omega^{<\omega}}$ be the 2^{k-1}-fold widening of $(b_\eta)_{\eta \in \omega^{<\omega}}$ at level $n+1$. Let $(\chi_i(x; w_i) : i < \omega)$ be defined as follows: if $i < n+1$, let $w_i = z_i$ and $\chi_i(x; w_i) = \psi_i(x; z_i)$. If $i \geq n+1$, let $w_i = (z_i^0, \ldots, z_i^{2^{k-1}-1})$ a tuple of variables consisting of 2^{k-1} copies of z_i. Then put

$$\chi_i(x; w_i) = \bigwedge_{j < 2^{k-1}} \psi_i(x; z_i^j).$$

By Lemma 1.7.3, $(c_\eta)_{\eta \in \omega^{<\omega}}$ is s-indiscernible and, by construction, $(\chi_i(x; w_i) : i < \omega)$, $(c_\eta)_{\eta \in \omega^{<\omega}}$ is a cdt-pattern and, moreover, if $i \neq j$

$$\{\chi_{n+1}(x; c_{0^n \frown \langle i \rangle}), \chi_{n+1}(x; c_{0^n \frown \langle j \rangle})\}$$

is inconsistent. For all $m < \omega$ and $\eta \in \omega^{<\omega}$, define $\varphi_m^{n+1} = \xi_m$ and $a_\eta^{n+1} = c_\eta$. We have satisfied requirements (1)-(3) and since our construction did not modify the formulas and parameters with level at most n, the construction never injures requirement (4).

Finally, define a cdt-pattern $(\varphi_n^\infty : n < \omega)$, $(a_\eta^\infty)_{\eta \in \omega^{<\omega}}$ by $\varphi_n^\infty = \varphi_n^n$ and $a_\eta^\infty = a_\eta^{l(\eta)}$. Our construction gives

7. $(a_\eta^\infty)_{\eta \in \omega^{<\omega}}$ is s-indiscernible.

8. If $\eta \in \omega^\omega$, $\{\varphi^\infty(x; a_{\eta|n}^\infty) : n < \omega\}$ is consistent.

9. If $\nu \lhd \eta$, then $\models (\forall x)[\varphi_{l(\eta)}^\infty(x; a_\eta^\infty) \to \varphi_{l(\nu)}^\infty(x; a_\nu^\infty)]$.

10. For all n, and $i \neq j$ $\{\varphi_{n+1}^\infty(x; a_{0^n \frown \langle i \rangle}^\infty), \varphi_{n+1}^\infty(x; a_{0^n \frown \langle j \rangle}^\infty)\}$ is inconsistent.

By s-indiscernibility, (9) and (10) imply that if $\eta \perp \nu$, then

$$\{\varphi_{l(\eta)}^\infty(x; a_\eta^\infty), \varphi_{l(\nu)}^\infty(x; a_\nu^\infty)\}$$

is inconsistent. This shows $(\varphi_n^\infty : n < \omega)$ and $(a_\eta^\infty)_{\eta \in \omega^{<\omega}}$ form an *sct*-pattern. We have thus shown $\kappa_{sct}^n \geq \aleph_1$. $\square$

We obtain the main theorem of this section.

Theorem 1.3.14. *If T is countable, then $\kappa_{cdt}(T) = \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$. Moreover, $\kappa_{\mathrm{cdt}}^n(T) = \kappa_{\mathrm{sct}}^n(T) + \kappa_{\mathrm{inp}}^n(T)$, provided $\kappa_{cdt}^n(T)$ is infinite.*

Proof. By Observation 1.3.5, $\kappa^n_{\text{cdt}}(T) \geq n$ for any T and $\kappa_{\text{cdt}}(T) > |T|^+$ if and only if $\kappa_{\text{cdt}}(T) = \infty$. It follows that, for countable theories, the possible values of $\kappa_{\text{cdt}}(T)$, and the only possible infinite values of $\kappa^n_{\text{cdt}}(T)$, are $\aleph_0$, $\aleph_1$, and ∞. The case of $\aleph_0$ is treated in Proposition 1.3.10, $\aleph_1$ is handled by Proposition 1.3.13, and for ∞ the result follows from Shelah's theorem (Fact 1.3.2). □

1.4 TP_1 and weak $k - \mathrm{TP}_1$

Say that a subset $\{\eta_i : i < k\} \subseteq \omega^{<\omega}$ is a collection of *distant siblings* if given $i \neq i'$, $j \neq j'$, all of which are $< k$, $\eta_i \wedge \eta_{i'} = \eta_j \wedge \eta_{j'}$.

Definition 1.4.1. Fix $k \geq 2$.

1. The formula $\varphi(x; y)$ has SOP_2 if there is a collection of tuples $(a_\eta)_{\eta \in 2^{<\omega}}$ satisfying the following.
   a) For all $\eta \in 2^\omega$, $\{\varphi(x; a_{\eta|\alpha}) : \alpha < \omega\}$ is consistent.
   b) If $\eta, \nu \in 2^{<\omega}$ and $\eta \perp \nu$, then $\{\varphi(x; a_\eta), \varphi(x; a_\nu)\}$ is inconsistent.
2. The formula $\varphi(x; y)$ has *weak k-TP*$_1$ if there is a collection of tuples $(a_\eta)_{\eta \in \omega^{<\omega}}$ satisfying the following.
   a) For all $\eta \in \omega^\omega$, $\{\varphi(x; a_{\eta|\alpha}) : \alpha < \omega\}$ is consistent.
   b) If $\{\eta_i : i < k\} \subseteq \omega^{<\omega}$ is a collection of distinct distant siblings, then $\{\varphi(x; a_{\eta_i}) : i < k\}$ is inconsistent.
3. The formula $\varphi(x; y)$ has *k-TP*$_1$ if there is a collection of tuples $(a_\eta)_{\eta \in \omega^{<\omega}}$ satisfying the following.
   a) For all $\eta \in \omega^\omega$, $\{\varphi(x; a_{\eta|\alpha}) : \alpha < \omega\}$ is consistent.
   b) If $\{\eta_i : i < k\} \subseteq \omega^{<\omega}$ is a collection of distinct pairwise incomparable nodes, then $\{\varphi(x; a_{\eta_i}) : i < k\}$ is inconsistent.
4. The theory T has either of the above properties if some formula does.

We remark that TP_1 is equivalent to SOP_2 in a strong way:

Fact 1.4.2. If a theory has TP_1 witnessed by a formula φ, then the theory also has SOP_2 witnessed by the same formula, and vice versa.

We recall the argument from [KK11]. Suppose $\varphi(x; y)$ witnesses SOP_2 with respect to the tree of parameters $(b_\eta)_{\eta \in 2^{<\omega}}$. Define a map $h : \omega^{<\omega} \to 2^{<\omega}$ recursively by $h(\emptyset) = \emptyset$ and $h(\beta \frown \langle i \rangle) = h(\beta) \frown 1^i \frown 0$, where 1^i denotes the all 1's sequence of length i. It is straightforward to check that $\varphi(x; y)$ witnesses TP_1 with respect to the parameters

$(b_{h(\eta)})_{\eta\in\omega^{<\omega}}$. The converse is obvious. Although SOP_2 and TP_1 are equivalent, it will be important for us to notationally distinguish them, as various combinatorial constructions are simplified by a judicious choice of the index set.

In [KK11], Kim and Kim show that k-TP_1 is equivalent to TP_1 for all $k \geq 2$, but the questions of whether weak k-TP_1 is equivalent to TP_1 was left unresolved. Using strongly indiscernible trees, we settle this, as well as show that TP_1 is always witnessed by a formula in a single free variable.

Finding and manipulating indiscernible witnesses

Lemma 1.4.3. *1. If T has weak k-TP_1 witnessed by $\varphi(x; y)$ then there is a strongly indiscernible tree $(a_\eta)_{\eta\in\omega^{<\omega}}$ witnessing this.*

2. If $\varphi(x; y)$ has TP_1 *then there is a strongly indiscernible tree witnessing this.*

3. If $\varphi(x, y)$ has SOP_2*, then there is a strongly indiscernible tree $(a_\eta)_{\eta\in 2^{<\omega}}$ witnessing this.*

Proof. (1) This was observed in [TT12], but we sketch a proof here for completeness. Let $(b_\eta)_{\eta\in\omega^{<\omega}}$ be a tree of tuples with respect to which $\varphi(x; y)$ witnesses weak k-TP_1. Let $(a_\eta)_{\eta\in\omega^{<\omega}}$ be locally based on the tree $(b_\eta)_{\eta\in\omega^{<\omega}}$. Suppose $\eta_0, \ldots, \eta_{n-1} \in \omega^{<\omega}$ lie along a path and let $\psi(y_0, \ldots, y_{n-1})$ denote the formula $(\exists x) \bigwedge_{i<n} \varphi(x; y_i)$. Then there are $\nu_0, \ldots, \nu_{n-1} \in \omega^{<\omega}$ so that

$$\mathrm{qftp}_{L_0}(\eta_0, \ldots, \eta_{n-1}) = \mathrm{qftp}_{L_0}(\nu_0, \ldots, \nu_{n-1})$$

and

$$\mathrm{tp}_\psi(a_{\eta_0}, \ldots, a_{\eta_{n-1}}) = \mathrm{tp}_\psi(b_{\nu_0}, \ldots, b_{\nu_{n-1}}).$$

The first equality implies that $\nu_0, \ldots, \nu_{n-1}$ all lie along a path so $\{\varphi(x; b_{\nu_i}) : i < n\}$ is consistent. By the second equality, $\{\varphi(x; a_{\eta_i}) : i < n\}$ is consistent. By compactness, this shows that all paths are consistent. Showing that any k distinct distant siblings remain inconsistent is similar. So $\varphi(x; y)$ witnesses weak k-TP_1 with respect to the tree $(a_\eta)_{\eta\in\omega^{<\omega}}$.

(2) This follows from (1) as weak 2-TP_1 and TP_1 are the same.

(3) By Fact 1.4.2, $\varphi(x, y)$ has TP_1. Now by (2), we may find a strongly indiscernible tree $(a_\eta)_{\eta\in\omega^{<\omega}}$ such that φ witnesses TP_1 with respect to $(a_\eta)_{\eta\in\omega^{<\omega}}$. Making the identification $2^{<\omega} = \{\eta \in \omega^{<\omega} : \eta(k) \in \{0, 1\} \text{ for all } k < l(\eta)\}$, it is easy to see that $(2^{<\omega}, \triangleleft, <_{lex}, \wedge)$ is an L_0-substructure of $(\omega^{<\omega}, \trianglelefteq, <_{lex}, \wedge)$ since $2^{<\omega}$ is closed under the $\wedge$-function and all the symbols in L_0 acquire their natural interpretation on $2^{<\omega}$ via restriction from $\omega^{<\omega}$. It follows that if $\eta_0, \ldots, \eta_{n-1}$ and $\nu_0, \ldots, \nu_{n-1}$ are two sequences from $2^{<\omega}$ with

$$\mathrm{qftp}_{L_0}(\eta_0, \ldots, \eta_{n-1}) = \mathrm{qftp}_{L_0}(\nu_0, \ldots, \nu_{n-1})$$

in $2^{<\omega}$, then this equality also holds in $\omega^{<\omega}$ and hence

$$\mathrm{tp}(a_{\eta_0}, \ldots, a_{\eta_{n-1}}) = \mathrm{tp}(a_{\nu_0}, \ldots, a_{\nu_{n-1}}),$$

so $(a_\eta)_{\eta\in 2^{<\omega}}$ is strongly indiscernible. Moreover, paths in 2^ω are paths also in ω^ω and incomparables in $2^{<\omega}$ remain incomparable when considered as elements in $\omega^{<\omega}$ so it is clear that $\varphi(x;y)$ will witness SOP_2 with respect to $(a_\eta)_{\eta\in 2^{<\omega}}$. □

Remark 1.4.4. We *aren't* making the (ostensibly) stronger claim that if $\varphi(x;y)$ witnesses SOP_2 with respect to the tree $(b_\eta)_{\eta\in 2^{<\omega}}$ then there is a strongly indiscernible tree $(a_\eta)_{\eta\in 2^{<\omega}}$ *based* on it — the proof of the existence of a strongly indiscernible tree witness involved going through TP_1 and then restricting.

Lemma 1.4.5. *1. If $(a_\eta)_{\eta\in\omega^{<\omega}}$ is a strongly indiscernible tree and $\varphi(x;y)$ is a formula so that for some $\eta\in\omega^\omega$, $\{\varphi(x;a_{\eta|n}) : n<\omega\}$ is consistent and for some $\xi\in\omega^{<\omega}$, $\{\varphi(x;a_{\xi\frown 0}),\varphi(x;a_{\xi\frown 1})\}$ is inconsistent, then T has* TP_1.

2. If $(a_\eta)_{\eta\in 2^{<\omega}}$ is a strongly indiscernible tree and $\varphi(x;y)$ is a formula so that for some $\eta\in 2^\omega$, $\{\varphi(x;a_{\eta|n}) : n<\omega\}$ is consistent and for some $\eta\in 2^{<\omega}$, $\{\varphi(x;a_{\eta\frown 0}),\varphi(x;a_{\eta\frown 1})\}$ is inconsistent, then T has SOP_2.

Proof. Both parts are immediate by Lemma 1.2.2, (1) and (2). □

Lemma 1.4.6. *(Path Collapse) Suppose κ is an infinite cardinal, $(a_\eta)_{\eta\in 2^{<\kappa}}$ is a tree strongly indiscernible over a set of parameters C and, moreover, $(a_{0^\alpha} : 0<\alpha<\omega)$ is indiscernible over cC. Let*

$$p(y;\bar{z}) = \mathrm{tp}(c;(a_{0\frown 0^\gamma} : \gamma<\kappa)/C).$$

Then if

$$p(y;(a_{0\frown 0^\gamma})_{\gamma<\kappa}) \cup p(y;(a_{1\frown 0^\gamma})_{\gamma<\kappa})$$

is not consistent, then T has SOP_2*, witnessed by a formula with free variables y.*

Proof. We may add C to the language, so assume $C=\emptyset$. With p defined as above, suppose

$$p(y;(a_{0\frown 0^\gamma} : \gamma<\kappa)) \cup p(y;(a_{1\frown 0^\gamma} : \gamma<\kappa))$$

is inconsistent. Then by indiscernibility and compactness, there is a formula ψ and $n<\omega$ so that

$$\{\psi(y;a_0,\ldots,a_{0\frown 0^{n-1}})\} \cup \{\psi(y;a_1,a_{10},\ldots,a_{1\frown 0^{n-1}})\}$$

is inconsistent. Let $(b_\eta)_{\eta\in 2^{<\kappa}}$ denote the n-fold elongation of $(a_\eta)_{\eta\in 2^{<\kappa}}$. By Lemma 1.2.7, $(b_\eta : \eta\in 2^{<\kappa})$ is strongly indiscernible. Since $c\models\{\psi(y;b_{0^\alpha}) : \alpha<\kappa\}$ and $\psi(y;b_0)\wedge\psi(y;b_1)$ is inconsistent (by strong indiscernibility), by Lemma 1.4.5, ψ witnesses SOP_2. □

Remark 1.4.7. It is significant that the type p does *not* contain $a_\emptyset$ as a parameter. As b_0 and b_1 are incomparable and $\psi(x;b_0)$ and $\psi(x;b_1)$ are inconsistent, we can conclude that $\psi(x;b_\eta)$ and $\psi(x;b_\nu)$ are inconsistent for all incomparable η,ν by strong indiscernibility. But, for example, strong indiscernibility does not guarantee $b_{0\frown 0}b_{0\frown 1}$ has the same type as b_0b_1 over $a_\emptyset$ as $0\wedge 1=\emptyset$ while $0^{n-1}\frown 0\wedge 0^{n-1}\frown 1 = 0^{n-1}$.

We now give two applications of the path-collapse lemma.

Weak $k - \mathrm{TP}_1$

Theorem 1.4.8. *Given $k \geq 2$, T has weak k-TP_1 if and only if T has TP_1.*

Proof. We will show that if T has weak k-TP_1, then T has SOP_2. Let $\varphi(x;y)$ witness weak k-TP_1 with respect to the strongly indiscernible tree $(a_\eta)_{\eta\in\omega^{<\omega}}$. Let n be maximal so that

$$\{\varphi(x; a_{\langle i\rangle\frown 0^\alpha}) : i < n, \alpha < \omega\}$$

is consistent. By definition of weak k-TP_1, n is at least 1 and at most $k-1$. Let $C = \{a_{\langle i\rangle\frown 0^\alpha} : i < n-1, \alpha < \omega\}$ (and put $C = \emptyset$ in the case that $n = 1$). Given $\eta \in \omega^{<\omega}$, let $\hat{\eta}$ be defined by

$$\hat{\eta}(i) = \begin{cases} \eta(i) + n - 1 & \text{if } i = 0 \\ \eta(i) & \text{otherwise,} \end{cases}$$

for all $i < l(\eta)$. The tree $(b_\eta)_{\eta\in\omega^{<\omega}}$ defined by $b_\eta = a_{\hat{\eta}}$ is strongly indiscernible over C. By choice of n,

$$\{\varphi(x; a_{\langle i\rangle\frown 0^\alpha}) : i < n, \alpha < \omega\}$$

is consistent, so let c realize it. By compactness, Ramsey, and automorphism, we may assume $(b_{0^\alpha} : 0 < \alpha < \omega)$ (i.e. $(a_{\langle n-1\rangle\frown 0^\alpha} : \alpha < \omega)$) is indiscernible over c. Letting the type p be defined by

$$p(y; \overline{z}) = \mathrm{tp}(c; (b_{0\frown 0^\alpha} : \alpha < \alpha)/C),$$

and unravelling definitions, we see that the type

$$p(y; (b_{0\frown 0^\alpha} : \alpha < \omega)) \cup p(y; (b_{1\frown 0^\alpha} : \alpha < \omega))$$

implies $\{\varphi(x; a_{\langle i\rangle\frown 0^\alpha}) : i < n+1, \alpha < \omega\}$ and is therefore inconsistent by the choice of n. By path-collapse, we've shown that T has SOP_2, completing one direction. The other direction is obvious. □

Reducing to one variable

Proposition 1.4.9. *Suppose T witnesses SOP_2 via $\varphi(x, y; z)$. Then there is a formula $\varphi_0(x; v)$ with free variables x and parameter variables v, or a formula $\varphi_1(y; w)$ with free variables y and parameter variables w so that one of φ_0 and φ_1 witness SOP_2.*

Proof. Let $\varphi(x, y; z)$ witness SOP_2 with respect to the strongly indiscernible tree $(a_\eta)_{\eta\in 2^{<\omega}}$. The first path is consistent and it is an indiscernible sequence so it follows that there is some $(c, c_0) \models \{\varphi(x, y; a_{0^\alpha}) : \alpha < \omega\}$ and such that moreover $(a_{0^\alpha} : \alpha < \omega)$ is indiscernible over c_0 (by Ramsey, automorphism, and compactness).

Define the function $h : 2^{<\omega} \to 2^{<\omega}$ recursively by $h(\emptyset) = \emptyset$ and $h(\eta \frown \langle i\rangle) = h(\eta) \frown 0 \frown \langle i\rangle$. Define the tree $(b_\eta)_{\eta\in 2^{<\omega}}$ by $b_\eta = a_{h(\eta)}$. It is proved in Lemma 1.7.7(1) that $(b_\eta)_{\eta\in 2^{<\omega}}$ is a strongly indiscernible tree. For each n, define a map $h_n : 2^{<\omega} \to 2^{<\omega}$ by

$$h_n(\eta) = \begin{cases} h(\eta) & \text{if } l(\eta) \leq n \\ h(\nu) \frown \xi & \text{if } \eta = \nu \frown \xi, l(\nu) = n. \end{cases}$$

By Lemma 1.7.7(2), the tree $(d_{n,\eta})_{\eta\in 2^{<\omega}}$ defined by $d_{n,\eta} = a_{h_n(\eta)}$ is strongly indiscernible as well. Moreover, as paths in $(b_\eta)_{\eta\in 2^{<\omega}}$ and $(d_{n,\eta})_{\eta\in 2^{<\omega}}$ are contained in paths in $(a_\eta)_{\eta\in 2^{<\omega}}$ and incomparable elements in these trees correspond to incomparable elements in $(a_\eta)_{\eta\in 2^{<\omega}}$, φ witnesses SOP_2 with respect to these trees of parameters as well.

Assume that no formula in the variable y has SOP_2. By induction, we will choose c_n so that

$$\{\varphi(x, c_n; d_{n,\eta|m}) : m < n\} \cup \{\varphi(x, c_n; d_{n,\eta\frown 0^\alpha}) : \alpha < \omega\} \tag{*}$$

is consistent for every $\eta \in 2^{\leq n}$.

For this, consider $(d^{(n)}_{n,\eta})_{\eta\in 2^{<\omega}}$, the nth-fattening of $(d_{n,\eta})$, and let $C_n = (d_{n,\eta} : \eta \in 2^{<n})$. By induction we show:

Claim. There is c_{n+1} such that $\left((d^{(n+1)}_{n+1,0^\alpha}) : \alpha < \omega\right)$ is indiscernible over $c_{n+1}C_n$ and

$$c_n\left(d^{(n)}_{n,0\frown 0\frown 0^\alpha}\right) \equiv_{d^{(n)}_{n,\emptyset}C_n} c_{n+1}\left(d^{(n)}_{n,0\frown 0\frown 0^\alpha}\right) \equiv_{d^{(n)}_{n,\emptyset}C_n} c_{n+1}\left(d^{(n)}_{n,0\frown 1\frown 0^\alpha}\right).$$

Note that $d^{(n)}_{n,\emptyset}C_n = C_{n+1}$.

Proof: The base case is above. Let

$$p_n(y, \overline{z}) = \mathrm{tp}\left(c_n, (d^{(n)}_{n,0\frown 0\frown 0^\alpha} : \alpha < \omega)/(d_{n,\emptyset})^{(n)}C_n\right).$$

By the path-collapse lemma,

$$p_n\left(y, \left((d^{(n)}_{n,0\frown 0\frown 0^\alpha}) : \alpha < \omega\right)\right) \cup p_n\left(y, \left((d^{(n)}_{n,0\frown 1\frown 0^\alpha}) : \alpha < \omega\right)\right)$$

is consistent. Let c_{n+1} realize it. Moreover, as

$$\left(d^{(n)}_{n,0\frown 0\frown 0^\alpha}, d^{(n)}_{n,0\frown 1\frown 0^\alpha}\right)_{\alpha<\omega} = \left(d^{(n+1)}_{n+1,0^\alpha}\right)_{\alpha<\omega}$$

is an indiscernible sequence, by Ramsey, automorphism, and compactness we may assume that it is indiscernible over $c_{n+1}C_n$. This shows (*).

By the definition of the trees $(d_{n,\eta})_{\eta\in 2^{<\omega}}$, we have shown that

$$\{\varphi(x, c_n; b_{\eta|m}) : m < n\} \cup \{\varphi(x, c_n; b_{\eta\frown 0^\alpha}) : \alpha < \omega\}$$

is consistent for each n and $\eta \in 2^{\leq n}$. By compactness, we can find one c which works for all possible paths in 2^ω simultaneously, giving a tree $(c, b_\eta)_{\eta\in 2^{<\omega}}$ witnessing SOP_2 for $\varphi(x; y, z)$. □

Remark 1.4.10. The necessity of defining the trees $(b_\eta)_{\eta\in 2^{<\omega}}$ and $(d_{n,\eta})_{\eta\in 2^{<\omega}}$ via h and h_n, respectively, stems from a technical obstacle in applying the path-collapse lemma: starting with the tree $(a_\eta)_{\eta\in 2^{<\omega}}$, we cannot apply the path collapse lemma directly to the type

$$q(y; (a_{0^\alpha} : \alpha < \omega)) = \mathrm{tp}(c_0/(a_{0^\alpha} : \alpha < \omega)),$$

as this type has $a_\emptyset$ as a parameter (see Remark 1.4.7 above). This is corrected by the offset functions h and h_n, allowing us to apply the path-collapse lemma 'higher' in the tree, where the parameters of interest are indiscernible over what we have constructed so far.

Corollary 1.4.11. *1. T has SOP_2 if and only if there is some formula in a single free variable witnessing this*

2. T has TP_1 if and only if there is some formula in a single free variable witnessing this

At this point it is natural to ask if $\kappa^1_{\mathrm{sct}} = \kappa^n_{\mathrm{sct}}$ holds for arbitrary n, at least for countable theories. Corollary 1.4.11 resolves the case of ∞, and we remark that the case of $\aleph_1$ follows from a standard argument in simplicity theory.

Proposition 1.4.12. *Any theory satisfies $\kappa^1_{\mathrm{cdt}} = \kappa^n_{\mathrm{cdt}}$, for all $n \in \omega$.*

Proof. The following are equivalent (see e.g. [Cas11, Proposition 3.8]).

1. $\kappa^n_{\mathrm{cdt}} \leq \kappa$.
2. For any type $p(x) \in S_n(A)$, there is some $A_0 \subseteq A$ such that $|A_0| < \kappa$ and p does not divide over A_0.

Clearly $\kappa^n_{\mathrm{cdt}} \geq \kappa^1_{\mathrm{cdt}}$. Assume now that $\kappa^1_{\mathrm{cdt}} \leq \kappa$ for some κ. We show by induction that (2) above holds for all n with respect to κ. Given $a_1 \dots a_n a_{n+1}$ and A, it follows by the inductive assumption that $a_1 \dots a_n \mathop{\smile\!\!\!\!\!\!\;|}_{A_0} A$ for some $A_0 \subseteq A$ with $|A_1| < \kappa$ and $a_{n+1} \mathop{\smile\!\!\!\!\!\!\;|}_{A_1 a_1 \dots a_n} A a_1 \dots a_n$ for some $A_1 \subseteq A$ with $|A_1| < \kappa$. Combined this implies (by left transitivity and right base monotonicity of dividing in arbitrary theories, see e.g. [CK12, Section 2]) that $a_1 \dots a_n a_{n+1} \mathop{\smile\!\!\!\!\!\!\;|}_{A_0 A_1} A$ and $|A_0 \cup A_1| < \kappa$. □

Corollary 1.4.13. *If $\kappa^n_{\mathrm{sct}} \geq \aleph_1$ then $\kappa^1_{\mathrm{sct}} \geq \aleph_1$.*

Proof. By Proposition 1.3.13, it is enough to show that $\kappa^1_{\mathrm{cdt}} \geq \aleph_1$, which follows by assumption and Proposition 1.4.12. □

The case of $\aleph_0$ appears to involve more complicated combinatorics and we leave it for future work.

1.5 Independence and amalgamation in NSOP_1 theories

We recall the definition of SOP_1 from [DS04]:

Definition 1.5.1. A formula $\varphi(x; y)$ exemplifies SOP_1 if and only if there are $(a_\eta)_{\eta \in 2^{<\omega}}$ so that

- For all $\eta \in 2^\omega$, $\{\varphi(x; a_{\eta|n}) : n < \omega\}$ is consistent,
- If $\eta \frown 0 \trianglelefteq \nu \in 2^{<\omega}$, then $\{\varphi(x; a_{\eta\frown 1}), \varphi(x; a_\nu)\}$ is inconsistent.

Given an array $(c_{i,j})_{i<\omega, j<2}$, write $\bar{c}_i = (c_{i,0}, c_{i,1})$ and $\bar{c}_{<i}$ for $(\bar{c}_j)_{j<i}$.

Lemma 1.5.2. *Suppose* $(c_{i,j})_{i<\omega,j<2}$ *is an array and* $\varphi(x;y)$ *is a formula over* C *with*

1. *For all* $i<\omega$, $c_{i,0} \equiv_{C\overline{c}_{<i}} c_{i,1}$;
2. $\{\varphi(x;c_{i,0}) : i<\omega\}$ *is consistent;*
3. $j \leq i \implies \{\varphi(x;c_{i,0}), \varphi(x;c_{j,1})\}$ *is inconsistent,*

then T *is* SOP_1.

Proof. For each n, define a subtree T_n of $2^{<\omega}$ by

$$T_n = \{\eta \frown 0^\alpha : \eta \in 2^{\leq n}, \alpha < \omega\} \cup \{\eta \frown 0^\alpha \frown 1 : \eta \in 2^{\leq n}, \alpha < \omega\}.$$

Let $P(T_n) \subseteq 2^\omega$ be the set of infinite branches of T_n. Namely,

$$P(T_n) = \{\eta \frown 0^\omega : \eta \in 2^{\leq n}\}.$$

As a first step, by induction on n we build an ascending sequence of trees $(l_\eta, r_\eta)_{\eta \in T_n}$, so that:

1. if $\eta \in P(T_n)$, $(l_{\eta|\alpha}, r_{\eta|\alpha})_{\alpha<\omega} \equiv_C (c_{\alpha,0}, c_{\alpha,1})_{\alpha<\omega}$,
2. if $\eta \frown 0 \in T_n$ then $r_{\eta\frown 0} = l_{\eta\frown 1}$,
3. if $\eta \in 2^{\leq n}$ then $(l_{\eta\frown 0}, r_{\eta\frown 0}) \equiv_{Cl_{\trianglelefteq\eta}r_{\trianglelefteq\eta}} (l_{\eta\frown 1}, r_{\eta\frown 1})$.

For the $n=0$ case, define $l_{0^\alpha} = c_{\alpha,0}$, $r_{0^\alpha} = c_{\alpha,1}$ and $l_{0^\alpha\frown 1} = r_{0^\alpha\frown 0}$ for all $\alpha<\omega$. For each $\alpha<\omega$, we can choose $\sigma_\alpha \in \mathrm{Aut}(\mathbb{M}/C\overline{c}_{<\alpha})$ such that $\sigma_\alpha(c_{\alpha,0}) = c_{\alpha,1}$. Let $r_{0^\alpha\frown 1} = \sigma_{\alpha+1}(c_{\alpha+1,1}) = \sigma_{\alpha+1}(r_{0^\alpha\frown 0})$. This defines $(l_\eta, r_\eta)_{\eta\in T_0}$ satisfying (1)-(3).

Now by induction suppose $(l_\eta, r_\eta)_{\eta\in T_n}$ has been defined. Suppose $\eta \in P(T_{n+1}) \setminus P(T_n)$. Then there is $\nu \in 2^{\leq n}$ so that $\eta = \nu \frown 1 \frown 0^\omega$. Then $\nu \frown 1 \in T_n$ and, by induction,

$$(l_{\nu\frown 0}, r_{\nu\frown 0}) \equiv_{Cl_{\trianglelefteq\nu}r_{\trianglelefteq\nu}} (l_{\nu\frown 1}, r_{\nu\frown 1})$$

and $r_{\nu\frown 0} = l_{\nu\frown 1}$. Choose an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}/Cl_{\trianglelefteq\nu}r_{\trianglelefteq\nu})$ such that $\sigma(l_{\nu\frown 0}, r_{\nu\frown 0}) = l_{\nu\frown 1}, r_{\nu\frown 1}$. Then define

$$(l_{\nu\frown 1\frown 0^\alpha}, r_{\nu\frown 1\frown 0^\alpha}) = \sigma(l_{\nu\frown 0\frown 0^\alpha}, r_{\nu\frown 0\frown 0^\alpha}) \text{ and}$$

$$(l_{\nu\frown 1\frown 0^\alpha\frown 1}, r_{\nu\frown 1\frown 0^\alpha\frown 1}) = \sigma(l_{\nu\frown 0\frown 0^\alpha\frown 1}, r_{\nu\frown 0\frown 0^\alpha\frown 1})$$

for all $\alpha<\omega$. This completes the construction of $(l_\eta, r_\eta)_{\eta\in T_{n+1}}$, properties (1)–(3) are satisfied because of the inductive assumption. We obtain $(l_\eta, r_\eta)_{\eta\in 2^{<\omega}}$ as the union over all n of $(l_\eta, r_\eta)_{\eta\in T_n}$.

Now we check that with respect to the parameters $(l_\eta)_{\eta\in 2^{<\omega}}$, φ witnesses SOP_1. Fix any path $\eta \in 2^\omega$, we have to check that $\{\varphi(x;l_{\eta|\alpha}) : \alpha<\omega\}$ is consistent. But given any

n, $l_{\trianglelefteq(\eta|n)} \subset T_n$ and by (1), $l_{\trianglelefteq(\eta|n)} \equiv_C (c_{\alpha,0})_{\alpha \leq n}$ hence $\{\varphi(x; l_{\eta|\alpha}) : \alpha \leq n\}$ is consistent, as $\{\varphi(x; c_{\alpha,0}) : \alpha \leq n\}$ is consistent, by hypothesis. Then $\{\varphi(x; l_{\eta|\alpha}) : \alpha < \omega\}$ is consistent by compactness.

Now fix $\eta \perp \nu \in 2^{<\omega}$ so that $(\eta \wedge \nu) \frown 0 \trianglelefteq \eta$ and $(\eta \wedge \nu) \frown 1 = \nu$. We must check $\{\varphi(x; l_\eta), \psi(x; l_\nu)\}$ is inconsistent. As $\nu = (\eta \wedge \nu) \frown 1$, we know that $l_\nu = l_{(\eta\wedge\nu)\frown 1} = r_{(\eta\wedge\nu)\frown 0}$ by (2). Let $\xi = (\eta \wedge \nu) \frown 0$. Then $\xi \trianglelefteq \eta$ and $l_\nu = r_\xi$ so it suffices to show $\{\varphi(x; l_\eta), \varphi(x; r_\xi)\}$ is inconsistent. Let $n = l(\eta)$ and $m = l(\xi)$. Then $m \leq n$ and by (1), we have $(l_\eta, r_\xi) \equiv_C (c_{n,0}, c_{m,1})$. By hypothesis, this implies $\{\varphi(x; l_\eta), \varphi(x; r_\xi)\}$ is inconsistent, so we finish. □

Definition 1.5.3. Suppose $\mathop{\smile\!\!\!\!\!\mid}$ is an $\mathrm{Aut}(\mathbb{M})$-invariant ternary relation on small subsets of $\mathbb{M}$.

1. We say $\mathop{\smile\!\!\!\!\!\mid}$ satisfies *weak independent amalgamation over models* if, given $M \models T$, $b_0c_0 \equiv_M b_1c_1$ satisfying $b_i \mathop{\smile\!\!\!\!\!\mid}_M c_i$ for $i = 0, 1$ and $c_0 \mathop{\smile\!\!\!\!\!\mid}_M c_1$, there is b satisfying $bc_0 \equiv_M bc_1 \equiv_M b_0c_0$.

2. We say $\mathop{\smile\!\!\!\!\!\mid}$ satisfies *independent amalgamation over models* if, given $M \models T$, $b_0 \equiv_M b_1$ satisfying $b_i \mathop{\smile\!\!\!\!\!\mid}_M c_i$ for $i = 0, 1$ and $c_0 \mathop{\smile\!\!\!\!\!\mid}_M c_1$, there is b satisfying $bc_0 \equiv_M b_0c_0$ and $bc_1 \equiv_M b_1c_1$.

3. We say $\mathop{\smile\!\!\!\!\!\mid}$ satisfies *stationarity over models* if: given $M \models T$, if $b_0 \equiv_M b_1$ and $b_0 \mathop{\smile\!\!\!\!\!\mid}_M c, b_1 \mathop{\smile\!\!\!\!\!\mid}_M c$ then $b_0 \equiv_{Mc} b_1$.

Definition 1.5.4. Suppose A, B, C are small subsets of the monster $\mathbb{M}$.

1. We say $A \mathop{\smile\!\!\!\!\!\mid}^i_C B$ if and only if $\mathrm{tp}(A/BC)$ can be extended to a global type Lascar-invariant over C. We denote its dual by $\mathop{\smile\!\!\!\!\!\mid}^{ci}$ - i.e. $A \mathop{\smile\!\!\!\!\!\mid}^i_C B$ holds if and only if $B \mathop{\smile\!\!\!\!\!\mid}^{ci}_C A$.

2. We say $A \mathop{\smile\!\!\!\!\!\mid}^u_C B$ if and only if $\mathrm{tp}(A/BC)$ is finitely satisfiable in C. We denote its dual by $\mathop{\smile\!\!\!\!\!\mid}^h$ - i.e. $A \mathop{\smile\!\!\!\!\!\mid}^h_C B$ if and only if $B \mathop{\smile\!\!\!\!\!\mid}^u_C A$.

Suppose $q(x)$ and $r(y)$ are global M-invariant types. Recall that the product $q(x) \otimes r(y) \in S_{xy}(\mathbb{M})$ is defined by $q(x) \otimes r(y) = \mathrm{tp}(ab/\mathbb{M})$ where $b \models r$ and $a \models q|_{\mathbb{M}b}$.

Proposition 1.5.5. *Fix a model $M \models T$. Suppose $c_1 \mathop{\smile\!\!\!\!\!\mid}^i_M c_0$, $c_j \mathop{\smile\!\!\!\!\!\mid}^i_M b_j$ for $j = 0, 1$ and $b_0c_0 \equiv_M b_1c_1$, but there is no b such that $bc_0 \equiv_M bc_1 \equiv_M b_0c_0$. Then T has* SOP_1.

Proof. Let $p(x; y) = \mathrm{tp}(b_0c_0/M)$. Our assumption entails that $p(x; c_0) \cup p(x; c_1)$ is inconsistent. By compactness, there is some $\varphi(x; y) \in p(x; y)$ so that $\{\varphi(x; c_0), \varphi(x; c_1)\}$ is inconsistent. Fix a global M-invariant type r so that $c_0 \models r|_{M_{b_0}}$ and a global M-invariant type q so that $c_1 \models q|_{M_{c_0}}$. Then $c_1c_0 \models (q \otimes r)|_M$. Let $(c^i_1, c^i_0)_{1 \leq i < \omega}$ be a Morley sequence in $(q \otimes r)|_{Mb_0c_0c_1}$ and put $(c^0_1, c^0_0) = (c_1, c_0)$.

First, we note that $b_0 \models \{\varphi(x; c_0^i) : i < \omega\}$ so *a fortiori* $\{\varphi(x; c_0^i) : i < \omega\}$ is consistent. Secondly, for any $N < \omega$, we have

$$(c_0^1 c_1^1) \dots (c_0^N c_1^N) \mathop{\mathop{\smile\!\!\!\!\!\!\,|}\limits^{i}}\limits_{M} c_0 c_1$$

so by M-invariance and the fact that $c_0 \equiv_M c_1$, we know that

$$c_0 \equiv_{M c_0^1 c_1^1 \dots c_0^N c_1^N} c_1$$

Next, as $c_1^1 \models q|_{Mc_0c_1}$, we have $c_1^1 \equiv_{Mc_0} c_1$ and therefore $\{\varphi(x; c_0), \varphi(x; c_1^1)\}$ is inconsistent. As $(c_1^i, c_0^i)_{i<\omega}$ is an M-indiscernible sequence, we've shown the following.

1. If $X \subseteq \omega$ and $j < k$ for all $k \in X$, then $\{\varphi(x; c_0^k) : k \in X\} \cup \{\varphi(x; c_i^j)\}$ is consistent for $i = 0, 1$.

2. If $X \subseteq \omega$ and $j < k$ for all $k \in X$, then, writing $\bar{c}_X$ for an enumeration of $\{c_0^k c_1^k : k \in X\}$, we have $c_0^j \equiv_{M\bar{c}_X} c_1^j$.

3. If $j \leq k$ then $\{\varphi(x; c_0^j), \varphi(x; c_1^k)\}$ is inconsistent.

Now by compactness (reversing the ordering on the sequence of pairs), we can find an array $(d_{i,j})_{i<\omega,j<2}$ such that the following holds.

1. For all $i < \omega$, $d_{i,0} \equiv_{M\bar{d}_{<i}} d_{i,1}$;

2. $\{\varphi(x; d_{i,0}) : i < \omega\}$ is consistent;

3. $j \leq i \implies \{\varphi(x; d_{i,0}), \varphi(x; d_{j,1})\}$ is inconsistent.

By Lemma 1.5.2, this implies T has SOP_1. $\square$

The following argument is an elaboration on [Che14, Proposition 6.20], which, in turn, was an elaboration on an argument of Kim [Kim01, Proposition 2.6].

Proposition 1.5.6. *Assume $\varphi(x; y)$ witnesses* SOP_1*. Then there are M, c_0, c_1, b_0, b_1 so that $c_0 \mathop{\smile\!\!\!\!\!\!\,|}\limits^{u}_{M} c_1$, $c_0 \mathop{\smile\!\!\!\!\!\!\,|}\limits^{u}_{M} b_0$, $c_1 \mathop{\smile\!\!\!\!\!\!\,|}\limits^{u}_{M} b_1$, $b_0c_0 \equiv_M b_1c_1$ and $\models \varphi(b_0, c_0) \wedge \varphi(b_1, c_1)$ but $\varphi(x; c_0) \wedge \varphi(x; c_1)$ is inconsistent.*

Proof. Suppose T has SOP_1 witnessed by φ. By compactness, we may assume that we have a tree of tuples $(a_\eta)_{\eta \in 2^{<\kappa}}$ for κ large enough ($\geq 2^{|T|}$ suffices) so that

- For all $\eta \in 2^\kappa$, $\{\varphi(x; a_{\eta|\alpha}) : \alpha < \kappa\}$ is consistent
- $\eta \frown 0 \trianglelefteq \nu \in 2^{<\kappa}$, then $\{\varphi(x; a_{\eta \frown 1}), \varphi(x; a_\nu)\}$ is inconsistent.

Fix a Skolemization T^{Sk} of T and in what follows, we'll work modulo this expanded theory. We will construct a sequence $(\eta_i, \nu_i)_{i<\omega}$ of elements of $2^{<\kappa}$ satisfying the following.

1. a_{ν_i}, a_{η_i} have the same type over $a_{\eta_{<i}}, a_{\nu_{<i}}$
2. If $i < j$ then $\eta_i \lhd \eta_j$ and $\eta_i \lhd \nu_j$.
3. $(\eta_i \wedge \nu_i) \frown 0 \trianglelefteq \eta_i$ and $(\eta_i \wedge \nu_i) \frown 1 = \nu_i$.

Given n, suppose $(\eta_i, \nu_i : i < n)$ have been chosen satisfying (1)-(3). Consider the sequence $(a_{\eta_{n-1} \frown 0^\alpha \frown 1} : \alpha < \kappa)$. As κ is large enough, there are $\alpha < \beta < \kappa$ so that $a_{\eta_{n-1} \frown 0^\alpha \frown 1}, a_{\eta_{n-1} \frown 0^\beta \frown 1}$ have the same type over $(a_{\eta_{<n}}, a_{\nu_{<n}})$. Let $\nu_n = \eta_{n-1} \frown 0^\alpha \frown 1$ and $\eta_n = \eta_{n-1} \frown 0^\beta \frown 1$. Now (1) and (2) are clearly satisfied, and, as $\alpha < \beta$, $(\eta_n \wedge \nu_n) = \eta_{n-1} \frown 0^\alpha$ so (3) follows. This completes the construction.

Now we claim that $(a_{\eta_i}, a_{\nu_i})_{i<\omega}$ satisfies:

4. $\{\varphi(x; a_{\eta_i}) : i < \omega\}$ is consistent,
5. a_{ν_i}, a_{η_i} have the same type over $a_{\nu_{<i}}, a_{\eta_{<i}}$,
6. $\{\varphi(x; a_{\nu_i}), \varphi(x; a_{\nu_j})\}$ is inconsistent for $i \neq j$.

Here (5) is immediate from our choice of the sequence and we get (4) since $i < j$ implies $\eta_i \lhd \eta_j$ and paths are consistent. To see (6), notice that if $i < j$ then as $\eta_i \lhd \nu_j$ and $\eta_i \perp \nu_i$, we have $(\nu_i \wedge \nu_j) = (\eta_i \wedge \nu_i)$ and hence $(\nu_i \wedge \nu_j) \frown 0 \trianglelefteq \nu_j$ and $\nu_i = (\nu_i \wedge \nu_j) \frown 1$ from which (6) follows, using SOP_1.

By compactness and Ramsey, we can find b and $(a_{\eta_i}, a_{\nu_i})_{i \leq \omega+1}$ indiscernible over b, satisfying (4)-(6), and such that $b \models \{\varphi(x; a_{\eta_i}) : i \leq \omega + 1\}$. Let $M = \mathrm{Sk}(a_{\eta_i}, a_{\nu_i})_{i<\omega}$. Then we have $a_{\eta_{\omega+1}} \mathop{\smile\!\!\!\!\!\!\mid}^u_M b$ and $a_{\nu_\omega} \mathop{\smile\!\!\!\!\!\!\mid}^u_M a_{\eta_{\omega+1}}$ by indiscernibility. As $a_{\nu_\omega}, a_{\eta_\omega}$ start an M-indiscernible sequence, there is $\sigma \in \mathrm{Aut}(\mathbb{M}/M)$ sending $a_{\eta_\omega} \mapsto a_{\nu_\omega}$. Let $b' = \sigma(b)$. Then $b' \equiv_M b$, $a_{\nu_\omega} \mathop{\smile\!\!\!\!\!\!\mid}^u_M b'$ (as $a_{\eta_\omega} \mathop{\smile\!\!\!\!\!\!\mid}^u_M b$ by indiscernibility) and $\models \varphi(b'; a_{\nu_\omega})$. But $\{\varphi(x; a_{\eta_{\omega+1}}), \varphi(x; a_{\nu_\omega})\}$ is inconsistent by (5) and (6). As φ is an L-formula, M is, in particular, an L-model and $\mathop{\smile\!\!\!\!\!\!\mid}^u$ in the sense of T^{Sk} implies $\mathop{\smile\!\!\!\!\!\!\mid}^u$ in the sense of T. $\square$

Theorem 1.5.7. *The following are equivalent.*

1. $\mathop{\smile\!\!\!\!\!\!\mid}^{ci}$ *satisfies weak independent amalgamation: given any* $M \models T$, $b_0 c_0 \equiv_M b_1 c_1$ *so that* $c_1 \mathop{\smile\!\!\!\!\!\!\mid}^i_M c_0$ *and* $c_j \mathop{\smile\!\!\!\!\!\!\mid}^i_M b_j$ *for* $j = 0, 1$, *there is* b *so that* $bc_0 \equiv_M bc_1 \equiv_M b_0 c_0$.
2. $\mathop{\smile\!\!\!\!\!\!\mid}^h$ *satisfies weak independent amalgamation: given any* $M \models T$, $b_0 c_0 \equiv_M b_1 c_1$ *so that* $c_1 \mathop{\smile\!\!\!\!\!\!\mid}^u_M c_0$ *and* $c_j \mathop{\smile\!\!\!\!\!\!\mid}^u_M b_j$ *for* $j = 0, 1$, *there is* b *so that* $bc_0 \equiv_M bc_1 \equiv_M b_0 c_0$.
3. T *is* $NSOP_1$.

Proof. (1) $\implies$ (2) is clear.

(2) $\implies$ (3) is Proposition 1.5.6.

(3) $\implies$ (1) is Proposition 1.5.5. $\square$

Proposition 1.5.8. *Assume there is an $Aut(\mathbb{M})$-invariant independence relation $\mathop{\smile\!\!\!\!\!\mid}$ on small subsets of the monster $\mathbb{M} \models T$ such that it satisfies the following properties, for an arbitrary $M \models T$ and arbitrary tuples from $\mathbb{M}$.*

1. *Strong finite character: if $a \not\mathop{\smile\!\!\!\!\!\mid}_M b$, then there is a formula $\varphi(x, b, m) \in tp(a/bM)$ such that for any $a' \models \varphi(x, b, m)$, $a' \not\mathop{\smile\!\!\!\!\!\mid}_M b$.*
2. *Existence over models: $M \models T$ implies $a \mathop{\smile\!\!\!\!\!\mid}_M M$ for any a.*
3. *Monotonicity: $aa' \mathop{\smile\!\!\!\!\!\mid}_M bb' \implies a \mathop{\smile\!\!\!\!\!\mid}_M b$.*
4. *Symmetry: $a \mathop{\smile\!\!\!\!\!\mid}_M b \iff b \mathop{\smile\!\!\!\!\!\mid}_M a$.*
5. *Independent amalgamation: $c_0 \mathop{\smile\!\!\!\!\!\mid}_M c_1$, $b_0 \mathop{\smile\!\!\!\!\!\mid}_M c_0$, $b_1 \mathop{\smile\!\!\!\!\!\mid}_M c_1$, $b_0 \equiv_M b_1$ implies there exists b with $b \equiv_{c_0 M} b_0$, $b \equiv_{c_1 M} b_1$.*

Then T is NSOP$_1$.

Proof. **Claim** Let $M \models T$, then $a \mathop{\smile\!\!\!\!\!\mid}^u_M b \implies a \mathop{\smile\!\!\!\!\!\mid}_M b$.

Proof of claim. If $a \not\mathop{\smile\!\!\!\!\!\mid}_M b$ then by strong finite character, there is some $\varphi(x; m, b) \in \mathrm{tp}(a/Mb)$ so that $a' \not\mathop{\smile\!\!\!\!\!\mid}_M b$ for any a' with $\models \varphi(a'; m, b)$. However, as $a \mathop{\smile\!\!\!\!\!\mid}^u_M b$, it follows that there is some $a' \in M$ such that $\models \varphi(a'; m, b)$. Then $b \not\mathop{\smile\!\!\!\!\!\mid}_M a'$ by symmetry and $b \not\mathop{\smile\!\!\!\!\!\mid}_M M$ by monotonicity, contradicting existence.

Now assume towards contradiction that T has SOP$_1$, and let $M, c_0, c_1, b_0, b_1, \varphi(x; y)$ as given in Proposition 1.5.6. By the claim and symmetry of $\mathop{\smile\!\!\!\!\!\mid}$ we have $c_0 \mathop{\smile\!\!\!\!\!\mid}_M c_1, b_0 \mathop{\smile\!\!\!\!\!\mid}_M c_0, b_1 \mathop{\smile\!\!\!\!\!\mid}_M c_1$. As $\mathop{\smile\!\!\!\!\!\mid}$ satisfies independent amalgamation over models, there is some $b \mathop{\smile\!\!\!\!\!\mid}_M c_0c_1$, $b \equiv_{c_0 M} b_0$, $b \equiv_{c_1 M} b_1$. This contradicts the inconsistency of $\{\varphi(x; c_0), \varphi(x; c_1)\}$.

1.6 Examples of NSOP$_1$ theories

Vector spaces with a generic bilinear form

Let L denote the language with two sorts V and K containing the language of abelian groups for variables from V, the language of rings for variables from K, a function $\cdot : K \times V \to V$, and a function $[\,] : V \times V \to K$. T_∞ is the model companion of the L-theory asserting that K is a field, V is a K-vector space of infinite dimension with the action of K given by $\cdot$, and $[\,]$ is a non-degenerate bilinear form on V. If $(K, V) \models T_\infty$ then K is an algebraically closed field.

The theory T_∞ was introduced by Nicolas Granger in [Gra99], who observed that its completions are not simple, but nonetheless have a notion of independence called Γ-non-forking satisfying essentially all properties of forking in stable theories, except local character.

Definition 1.6.1. We are using the notation from [Gra99, Notation 9.2.4]. Let $M = (V, \tilde{K})$ be a sufficiently saturated model of T_∞. Let $A \subseteq B \subset M$ and $c \in M$ with c a singleton.

Let $c \mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B$ be the assertion that $K_{Ac} \mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}_{K_A} K_B$ in the sense of non-forking independence for algebraically closed fields and one of the following holds:

1. $c \in \tilde{K}$
2. $c \in \langle A \rangle$
3. $c \notin \langle B \rangle$ and $[c, B]$ is Φ-independent over A,

where '$[c, B]$ is Φ-independent over A' means that whenever $\{b_0, \ldots, b_{n-1}\}$ is a linearly independent set in $B_V \cap (V \setminus \langle A \rangle)$ then the set $\{[c, b_0], \ldots, [c, b_{n-1}]\}$ is algebraically independent over the field $K_B(K_{Ac})$.

By induction, for $c = (c_0, \ldots, c_m)$ define $c \mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B$ by

$$c \mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B \iff (c_0, \ldots, c_{m-1}) \mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B \text{ and } c_m \mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{Ac_0 \ldots c_{m-1}} Bc_0 \ldots c_{m-1}.$$

Fact 1.6.2. [Gra99, Theorem 12.2.2] Let $M = (V, K) \models T_\infty$. Then the relation on subsets of M given by Γ-non-forking is automorphism invariant, symmetric, and transitive. Moreover, it satisfies extension, finite character, and stationarity over a model.

Lemma 1.6.3. *If c is a tuple and A, B are small sets with $c \not\mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B$, then there is a formula $\varphi(x; a, b) \in tp(c/AB)$ so that*

$$\models \varphi(c'; a, b) \implies c' \not\mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B.$$

Proof. Suppose $c = (c_0, \ldots, c_{n-1})$ a tuple and $c \not\mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B$. Let k be maximal so that $(c_0, \ldots, c_{k-1}) \mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{A} B$. It follows that $c_k \not\mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\Gamma}_{Ac_0 \ldots c_{k-1}} Bc_0 \ldots c_{k-1}$, so one of the following possibilities occurs:

1. $K_{Ac_0 \ldots c_k} \not\mathop{\smile\hskip-0.9em\raise0.25ex\hbox{$\mid$}}^{\mathrm{ACF}}_{K_{Ac_0 \ldots c_{k-1}}} K_{Bc_0 \ldots c_{k-1}}$
2. $c_k \in \langle Bc_0 \ldots c_{k-1} \rangle \setminus \langle Ac_0 \ldots c_{k-1} \rangle$
3. There is a linearly independent set $\{d_0, \ldots, d_{l-1}\}$ from $(Bc_0 \ldots c_{k-1})_V \cap (V \setminus \langle Ac_0 \ldots c_{k-1} \rangle)$ so that $\{[c_k, d_0], \ldots, [c_k, d_{l-1}]\}$ is not algebraically independent over $K_{Bc_0 \ldots c_{k-1}}(K_{Ac_0 \ldots c_k})$.

The existence of the desired formula requires an argument only in case (3). In this case, there is a nonzero polynomial $p(x_0, \ldots, x_{l-1}; a, b, c_0, \ldots, c_k)$ with coefficients in $K_{Bc_0 \ldots c_{k-1}}(K_{Ac_0 \ldots c_k})$ so that $p([c_k, d_0], \ldots, [c_k, d_{l-1}]; a, b, c_0, \ldots, c_k) = 0$. By reindexing the d_j, we may assume that there is $m \leq l$ so that $d_j = c_{i_j}$ for $j < m$ and $d_j \in B$ for $j \geq m$. Let $d = (d_m, \ldots, d_{l-1})$. Writing $y = (y_0, \ldots, y_k)$, let $\chi(y; a, b, d)$ be the formula which asserts the following:

1. the polynomial $p(x_0, \ldots, x_{l-1}; a, b, y)$ is a nonzero polynomial;

2. the set $\{y_{i_0}, \ldots, y_{i_{m-1}}\} \cup \{d_m, \ldots, d_{l-1}\}$ is linearly independent;
3. $p([y_k, y_{i_0}], \ldots, [y_k, y_{i_{m-1}}], [y_k, d_m], \ldots, [y_k, d_{l-1}]; a, b, y) = 0$

Then $\chi(y; a, b, d) \in \mathrm{tp}(c/B)$ and if $\models \chi(c'; a, b, d)$ then it is easy to check $c' \not\mathop{\smile\hskip-0.9em\mid}^{\Gamma}_A B$. $\square$

Corollary 1.6.4. *The two-sorted theory T_∞ of infinite dimensional vector spaces over algebraically closed fields with a generic bilinear form is NSOP$_1$.*

ω-free PAC fields of characteristic zero

Definition 1.6.5. A field F is called *pseudo-algebraically closed* if every absolutely irreducible variety defined over F has an F-rational point. A field F is called *ω-free* if it has a countable elementary substructure F_0 with $\mathcal{G}(F_0) \cong \hat{\mathbb{F}}_\omega$, the free profinite group on countably many generators.

In [Cha99], Chatzidakis showed that a PAC field has a simple theory if and only if it has finitely many degree n extensions for all n so an ω-free PAC field will not be simple. Nonetheless, she showed that an ω-free PAC field comes equipped with a notion of independence which is well-behaved.

Fact 1.6.6. [Cha02] Suppose F is a sufficiently saturated ω-free PAC field of characteristic zero. Given $A = \mathrm{acl}(A)$, $B = \mathrm{acl}(B)$, $C = \mathrm{acl}(C)$ with $C \subseteq A, B \subseteq F$, write $A \mathop{\smile\hskip-0.9em\mid}^{I}_C B$ to indicate that $A \mathop{\smile\hskip-0.9em\mid}^{\mathrm{ACF}}_C B$ and $A^{\mathrm{alg}} B^{\mathrm{alg}} \cap \mathrm{acl}(AB) = AB$. Extend this to non-algebraically closed sets by stipulating $a \mathop{\smile\hskip-0.9em\mid}^{I}_D b$ holds if and only if $\mathrm{acl}(aD) \mathop{\smile\hskip-0.9em\mid}^{I}_{\mathrm{acl}(D)} \mathrm{acl}(bD)$. Then $\mathop{\smile\hskip-0.9em\mid}^{I}$ satisfies existence over models, monotonicity, symmetry, and independent amalgamation over models.

It remains to check that $\mathop{\smile\hskip-0.9em\mid}^{I}$ satisfies strong finite character. The proof of it was pointed out to us by Zoé Chatzidakis, whom we would like to thank.

Lemma 1.6.7. *Suppose F is a sufficiently saturated ω-free PAC field of characteristic zero. If a, b, c are tuples from F and $a \mathop{\smile\hskip-0.9em\mid}^{I}_c b$ then there is a formula $\varphi(x; b, c) \in tp(a/bc)$ so that if $F \models \varphi(a'; b, c)$ then $a' \not\mathop{\smile\hskip-0.9em\mid}^{I}_c b$.*

Proof. If $a \not\mathop{\smile\hskip-0.9em\mid}^{\mathrm{ACF}}_c b$, then the existence of such a formula is clear, so we may assume $a \mathop{\smile\hskip-0.9em\mid}^{\mathrm{ACF}}_c b$. As $a \not\mathop{\smile\hskip-0.9em\mid}^{I}_c b$, there are $\beta \in \langle cb \rangle^{\mathrm{alg}}$, $\alpha \in \langle ca \rangle^{\mathrm{alg}}$ not in F such that $F(\alpha) = F(\beta)$ and $\beta \notin F\langle c \rangle^{\mathrm{alg}}$. We choose them so that $F(\beta)$ is Galois over F (always possible since $F \cap \langle ca \rangle^{\mathrm{alg}} \langle cb \rangle^{\mathrm{alg}}$ is Galois over $(F \cap \langle ca \rangle^{\mathrm{alg}})(F \cap \langle cb \rangle^{\mathrm{alg}}) = \mathrm{acl}(ca)\,\mathrm{acl}(cb)$).

Some of the conjugates of β over $\langle cb \rangle$ might lie in $F\langle c \rangle^{\mathrm{alg}}$ and this will be witnessed by elements of $\mathrm{acl}(cb) = F \cap \langle cb \rangle^{\mathrm{alg}}$. We choose an element b' of $\mathrm{acl}(cb)$ such that $\langle cbb' \rangle$ contains $\langle cb\beta \rangle \cap F$ and $\langle cbb' \rangle$ is closed under $\mathrm{Aut}(\mathrm{acl}(cb)/\langle cb \rangle)$. Let the formula $\theta(y; b, c)$ isolate $\mathrm{tp}(b'/bc)$.

Let $P(Y,b,c)$ be a minimal polynomial of b' over $\langle bc\rangle$, and let $Q(Z,Y,b,c)$ be such that $Q(Z,b',b,c)$ is a minimal polynomial of β over $\langle cbb'\rangle$.

Claim. If $\models \theta(b_1,b,c)$, then $P(b_1,b,c)=0$, $Q(Z,b_1,b,c)$ is irreducible of degree $[\langle cb\beta\rangle : \langle cbb'\rangle]$ and a solution of Q defines a Galois extension, which is not contained in $F\langle c\rangle^{\text{alg}}$.

The first two assertions of the claim are immediate. For the last one, assume that (b_1,b_2) satisfies $P(b_1,b,c)=0 \wedge Q(b_2,b_1,b,c)=0$, and that $Q(Z,b_1,b,c)$ is irreducible and defines a Galois extension of the right degree (all this is expressible in $\text{tp}_F(b'/bc)$), but that $b_2 \in F\langle c\rangle^{\text{alg}}$. Then there is a formula in $\text{tp}_F(b_1/cb)$ which will say that such a b_2 exists, and is therefore not in $\text{tp}_F(b'/bc)$.

Similarly let $a' \in \text{acl}(ac)$ be such that $\langle ca\alpha\rangle \cap F = \langle caa'\rangle$ and let $R(W,T,c)$ be such that $R(W,a,c)$ is a minimal polynomial of a' over $\langle ca\rangle$ and let $S(X,W,T,c)$ be such that $S(X,a',a,c)$ is a minimal polynomial of α over $\langle caa'\rangle$.

The formula $\varphi(t,b,c)$ is a conjunction of the following assertions:

- $\exists y\theta(y,b,c)$,
- $R(W,t,c)$ is not the trivial polynomial,
- $(\exists w)R(w,t,c)=0$ and $S(X,w,t,c)$ is irreducible over F of degree $[\langle ca\alpha\rangle : \langle caa'\rangle]$,
- $(\forall z)[Q(z,y,b,c)=0 \to$ "$F(z)$ contains a root of $S(X,w,t,c)=0$".

These statements are first-order using standard facts on interpretability of finite algebraic extensions of a field in a field and definability of irreducibility.

Assume now that d satisfies $\varphi(t,b,c)$. Let $y=b_1$ and $w=d_1\in F$ be as guaranteed to exist by φ, and let b_2 be a root of $Q(Z,b_1,b,c)=0$; then $F(b_2)$ is a proper Galois extension of F of degree $[\langle cb\beta\rangle : \langle cbb'\rangle]$ which is not contained in $F\langle c\rangle^{\text{alg}}$.

Because d satisfies φ, if d_2 satisfies $S(X,d_1,d,c)=0$, then $F(d_2)=F(b_2)$. As $F(b_2)\not\subseteq F\langle c\rangle^{\text{alg}}$, we necessarily have $d\notin\langle c\rangle^{\text{alg}}$ and, therefore, either $d \not\mathop{\smile\!\!\!\!\!\!\mid}\,_c^{\text{ACF}} b$ or, otherwise, $\langle cd\rangle^{\text{alg}}\langle cb\rangle^{\text{alg}}\cap F \neq \text{acl}(cd)\text{acl}(cb)$. This shows $d \not\mathop{\smile\!\!\!\!\!\!\mid}\,_c^{I} b$. $\square$

Corollary 1.6.8. *The theory of ω-free PAC fields of characteristic 0 is* NSOP_1.

Examples via Parametrization

In this subsection, we show how to construct NSOP_1 theories from simple ones. We start with a simple theory T obtained as the theory of a Fraïssé limit satisfying the strong amalgamation property and, by analogy with the theory of parametrized equivalence relations T^*_{feq}, form the parametrization of this structure. We show that the resulting theories are NSOP_1 by proving an independence theorem for a natural independence notion associated to these theories. The construction we perform here was studied by Baudisch [Bau02] in the context of arbitrary model complete theories eliminating $\exists^\infty$. We expect that our results hold in this greater generality as well, but our setting already encompasses many interesting examples and simplifies the study of amalgamation.

We begin by recalling some facts from Fraïssé theory.

Definition 1.6.9. (SAP) Suppose $\mathbb{K}$ is a class of finite structures. We say $\mathbb{K}$ has the Strong Amalgamation Property (SAP) if given $A, B, C \in \mathbb{K}$ and embeddings $e : A \to B$ and $f : A \to C$ there is a $D \in \mathbb{K}$ and embeddings $g : B \to D$ and $h : C \to D$ so that the following diagram commutes:

$$\begin{tikzcd} & B \arrow[dr, "g"] & \\ A \arrow[ur, "e"] \arrow[dr, "f"'] & & D \\ & C \arrow[ur, "h"'] & \end{tikzcd}$$

and, moreover, $(\mathrm{im}g) \cap (\mathrm{im}h) = \mathrm{im}ge$ (and hence $= \mathrm{im}hf$, as well).

The following is a useful criterion for SAP:

Fact 1.6.10. [Hod93] Suppose $\mathbb{K}$ is the age of a countable ultrahomogeneous structure M. Then the following are equivalent:

1. $\mathbb{K}$ has the strong amalgamation property.

2. M has no algebraicity.

Let $\mathbb{K}$ denote a Fraïssé class in a finite relational language $L = \langle R_i : i < k\rangle$ where each relation symbol R_i has arity n_i. Let T the complete L-theory of the Fraïssé limit of $\mathbb{K}$. We'll define a new language L_{pfc} which contains two sorts P and O. For each $i < k$, there is an $(n_i + 1)$-ary relation symbol R^i_x where x is a variable of sort P and the suppressed n_i variables belong to the sort O.

Given an L_{pfc}-structure M, it is convenient to write $M = (A, B)$ where $O(M) = A$ and $P(M) = B$. We will refer to elements named by O as *objects* and elements named by P as *parameters*. Given $b \in B$, we define the *L-structure associated to b in M*, denoted A_b, to be the L-structure interpreted in M with domain A and each relation symbol R_i interpreted by $R^i_b(A)$. If $b \in B$ and $C \subseteq A$, write $\langle C\rangle_b$ to denote the L-substructure of A_b generated by C (as we assume the language is relational, this will have C as its domain).

We describe a class of finite structures $\mathbb{K}_{\mathrm{pfc}}$ to be the class defined in the following way. Let

$$\mathbb{K}_{\mathrm{pfc}} = \{M = (A, B) \in \mathrm{Mod}(L_{\mathrm{pfc}}) : |M| < \aleph_0, (\forall b \in B)(\exists D \in \mathbb{K})\,(A_b \cong D)\}$$

From now on, we'll assume $\mathbb{K}$ also satisfies SAP.

Lemma 1.6.11. *$\mathbb{K}_{\mathrm{pfc}}$ is a Fraïssé class satisfying the Strong Amalgamation Property (SAP).*

Proof. HP is clear and, as we allow the empty structure to be a model in $\mathbb{K}_{\text{pfc}}$, JEP follows from SAP. So we show SAP.

First, we may assume that 3 models in the amalgamation diagram have the same set of parameters. Suppose (A, D), (B, E) and (C, F) are in $\mathbb{K}_{\text{pfc}}$ and we have embeddings

$$\begin{tikzcd} & (C,F) \\ (A,D) \arrow[ur, "i"] \arrow[dr, "j"'] & \\ & (B,E) \end{tikzcd}$$

By moving F and E over D if necessary, we may assume that i and j are just the inclusion maps on parameters and that $F \cap E = D$. By SAP in $\mathbb{K}$, for each $d \in D$, there are embeddings f_d, g_d and $G_d \in \mathbb{K}$ so that the following diagram commutes,

$$\begin{tikzcd} & C_d \arrow[dr, "f_d"] & \\ A_d \arrow[ur, "i"] \arrow[dr, "j"'] & & G_d \\ & B_d \arrow[ur, "g_d"'] & \end{tikzcd}$$

where i and j are the induced maps, so that $f_d(C_d) \cap g_d(B_b) = (f_d \circ i)(A_d)$. Since the language is relational, HP implies that we may take $G_d = f_d(C_d) \cup g_d(D_d)$. Moreover, we may choose f_d and g_d so that they are the same functions for all $d \in D$ on the underlying sets C and B respectively. Call these functions f and g. Let G be the underlying set of G_d for some (all) $d \in B$. Now define a structure $(G, E \cup F)$ so that for all $d \in D = E \cap F$, G_d is as above, if $a \in E \setminus F$, G_a is some structure in $\mathbb{K}$ extending $g(B_a)$ and, likewise, if $a \in F \setminus E$, G_a is some structure extending $f(C_a)$. The functions f and g extend to embeddings $f : (C, F) \to (G, E \cup F)$ and $g : (B, E) \to (G, E \cup F)$ so that f and g are both inclusions on parameters. By construction, it is clear that $fi = gj$. Moreover, $fi(A) = f(C) \cap g(B)$ and $fi(D) = f(E) \cap g(F)$, which establishes SAP in $\mathbb{K}_{\text{pfc}}$. □

As $\mathbb{K}_{\text{pfc}}$ is a Fraïssé class, there is a unique countable ultrahomogeneous L_{pfc}-structure with age $\mathbb{K}_{\text{pfc}}$. Let T_{pfc} denote its theory. By Fraïssé theory, this theory eliminates quantifiers and is $\aleph_0$-categorical.

Lemma 1.6.12. *Suppose $(A, B) \models T_{\text{pfc}}$. Then, for all $b \in B$, $A_b \models T$.*

Proof. Since the property that for all $b \in B$, $A_b \models T$ is an elementary property, it suffices to check this when (A, B) is the unique countable model of T_{pfc}. If $d, e \in A_b$ satisfy $\text{tp}_L(d) = \text{tp}_L(e)$ then, by quantifier-elimination, it is easy to check $\text{tp}_{L_{\text{pfc}}}(b, d) = \text{tp}_{L_{\text{pfc}}}(b, e)$

and ultrahomogeneity of (A, B) implies there is an L_{pfc}-automorphism of (A, B) fixing b and taking d to e. The induced L-automorphism of A_b witnesses that A_b is ultrahomogeneous. By Fraïssé theory there is up to isomorphism a unique countable ultrahomogeneous L-structure with age $\mathbb{K}$ so A_b is isomorphic to a model of T, so $A_b \models T$. $\square$

Suppose $\mathbb{M} = (\mathbb{A}, \mathbb{B})$ is a monster model of T_{pfc}. Given a formula $\varphi \in L$ and a parameter $p \in \mathbb{B}$, define $\varphi_p \in L_{\text{pfc}}$ to be the formula obtained by replacing each occurrence of R_i by R^i_p and giving the objects their eponymous interpretations in $\mathbb{A}_p$ – formally, this defines φ_p for atomic φ and then the full definition follows by induction on the complexity of the formulas. If $C \subseteq \mathbb{A}$ is a set of objects and q is an L-type over C (considered as a subset of $\mathbb{A}_p$), we define the type q_p by

$$q_p = \{\varphi_p : \varphi \in q\}.$$

Lemma 1.6.13. *Suppose $\{p_i : i < \alpha\} \subseteq \mathbb{B}$ is a collection of distinct parameters and $q^i : i < \alpha)$ is a sequence of non-algebraic L-types over $C \subseteq \mathbb{A}$ (possibly with repetition), where q^i is considered as a type in $\mathbb{A}_{p_i}$. Then the L_{pfc}-type $\bigcup_{i<\alpha} q^i_{p_i}$ is consistent.*

Proof. By compactness, it suffices to consider the case where $\alpha < \omega$ and when the q^i are all finite types. Hence, we simply have to show

$$\mathbb{M} \models (\exists x) \bigwedge_{i<\alpha} q^i_{p_i}(x).$$

Moreover, by quantifier-elimination in T, we may assume that each q^i is quantifier-free. For each $i < \alpha$, let $C_i \in \text{Age}(\mathbb{A}_{p_i})$ the finite substructure generated by the elements of C mentioned in all of the q^i. So, the underlying set of each C_i is the same, although the interpretations of the relations may differ. Given any $i < \alpha$, we know that

$$\mathbb{A}_{p_i} \models (\exists x) \bigwedge q^i_{p_i}(x)$$

so there is $D_i \in \text{Age}(\mathbb{A}_{p_i})$ containing a witness d_i to the above existential formula. By non-algebraicity of each type, we may assume that $d_i \notin C_i$ and, by HP, that $D_i = C \cup \{d_i\}$.

Now define an L_{pfc}-structure E with underlying set of objects $C \cup \{*\}$ where $*$ is some new element and its parameters are $\{p_i : i < \alpha\}$, and the relations are interpreted so that for each $i < \alpha$, the map is the identity on C and sends $d_i \mapsto *$ is an isomorphism of L-structures from D_i to E_{p_i}. It is clear that $E \in \mathbb{K}_{\text{pfc}}$ so there is a copy F isomorphic over $C \cup \{p_i : i < \alpha\}$ to it in $\text{Age}(\mathbb{M})$. Now

$$F \models (\exists x) \bigwedge_{i<\alpha} q^i_{p_i}(x)$$

and hence this is satisfied in $\mathbb{M}$, so we're done. $\square$

Lemma 1.6.14. *Suppose $A, B, C \subseteq \mathbb{A}$ are small sets of objects, $F \subseteq \mathbb{B}$ is a small set of parameters, $A \cap B \subseteq C$, and $b_0, b_1 \in \mathbb{B}$ satisfy $b_0 \equiv_{CF} b_1$. Then there is some $b \in \mathbb{B}$ so that $b \equiv_{ACF} b_0$ and $b \equiv_{BCF} b_1$ (all in L_{pfc}).*

Proof. Given a set $D \subseteq \mathbb{A}$ and $p \in \mathbb{B}$, recall that we write $\langle D \rangle_p$ for the L-substructure of $\mathbb{A}_p$ with underlying set D. By compactness, it suffices to prove the lemma when A, B, C, and F are finite. By quantifier-elimination, demanding some $b \in \mathbb{B}$ so that $b \equiv_{AC} b_0$ and $b \equiv_{BC} b_1$ is equivalent to asking that $\langle AC \rangle_b \cong \langle AC \rangle_{b_0}$ and $\langle BC \rangle_b \cong \langle AC \rangle_{b_1}$. Now, as $b_0 \equiv_C b_1$, $\langle C \rangle_{b_0}$ may be identified with $\langle C \rangle_{b_1}$. We may view $C, \langle AC \rangle_{b_0}$, and $\langle BC \rangle_{b_1}$ as elements of $\mathbb{K}$. In $\mathbb{K}$, we have inclusions $i : C \to \langle AC \rangle_{b_0}$ and $j : C \to \langle BC \rangle_{b_1}$, so by SAP, there are embeddings f, g and a $D \in \mathbb{K}$ so that the following diagram commutes

$$\begin{tikzcd}
 & \langle AC \rangle_{b_0} \arrow[dr, "f"] & \\
C \arrow[ur, "i"] \arrow[dr, "j"'] & & D \\
 & \langle BC \rangle_{b_1} \arrow[ur, "g"'] &
\end{tikzcd}$$

where $f(AC) \cap g(BC) = C$. By HP, D may be taken to have $f(AC) \cup g(BC)$ as its domain. Since $A \cap B \subseteq C$, D is isomorphic over C to an L-structure with underlying set $A \cup B \cup C$, so we may assume that f and g are both inclusions. Let b_* denote some new parameter element outside of F and define a structure with parameter set $\{b_*, b_0, b_1\} \cup F$ and $A \cup B \cup C$ as its set of objects so that $\langle ABC \rangle_{b_*} \cong D$. This clearly defines a structure in $\mathbb{K}_{\text{pfc}}$. In the substructure with only $A \cup C$ as the set of objects, there is an automorphism fixing F taking b_* to b_0. This shows that $b_* \equiv_{ACF} b_0$ and a symmetric argument shows $b_* \equiv_{BCF} b_1$. It follows that we can find such a b_* in $\mathbb{B}$. $\square$

Towards proving an independence theorem for T_{pfc}, we will define a notion of independence for parameterized structures.

Definition 1.6.15. ($\mathop{\smile\!\!\!\!\!\mid}^{\text{pfc}}$)

1. Suppose $p \in \mathbb{B}$ is a parameter. Suppose $A, B, C \subseteq \mathbb{A}$. We define $\mathop{\smile\!\!\!\!\!\mid}^{p}$ by

$$A \mathop{\smile\!\!\!\!\!\mid}_{C}^{p} B \text{ in } \mathbb{M} \iff A \mathop{\smile\!\!\!\!\!\mid}_{C} B \text{ in } \mathbb{A}_p,$$

   where the undecorated $\mathop{\smile\!\!\!\!\!\mid}$ on the right-hand side denotes the usual non-forking independence – i.e. $\text{tp}(A/BC)$ does not fork over C.

2. If $A, B, C \subseteq \mathbb{A}$ and $D, E, F \subseteq \mathbb{B}$, we define $\mathop{\smile\!\!\!\!\!\mid}^{\text{pfc}}$ by

$$A, D \mathop{\smile\!\!\!\!\!\mid}_{C,F}^{\text{pfc}} B, E \iff D \cap E \subseteq F, \text{ and for all } p \in F, A \mathop{\smile\!\!\!\!\!\mid}_{C}^{p} B.$$

Proposition 1.6.16. *Assume T is a simple theory. Suppose $A, B \subseteq \mathbb{A}$ are small sets of objects and $D, E \subseteq \mathbb{B}$ are small sets of parameters and $M = (C, F)$ is a small model of T_{pfc} satisfying*

$$A, D \mathop{\smile\!\!\!\!\!\mid}_{C,F}^{\mathrm{pfc}} B, E$$

Suppose moreover that a_0, a_1 are tuples from $\mathbb{A}$ and b_0, b_1 are tuples from $\mathbb{B}$ satisfying $a_0, b_0 \mathop{\smile\!\!\!\!\!\mid}_{CF}^{\mathrm{pfc}} A, D$, $a_1, b_1 \mathop{\smile\!\!\!\!\!\mid}_{C,F}^{\mathrm{pfc}} B, E$ and $a_0, b_0 \equiv_{CF} a_1, b_1$. Then there are a from $\mathbb{A}$ and b from $\mathbb{B}$ so that $a, b \equiv_{ACDF} a_0, b_0$ and $a, b \equiv_{BCEF} a_1, b_1$.

Proof. First, we solve the amalgamation problem for objects. Without loss of generality, D, E, F are pairwise disjoint. By Lemma 1.6.12, we know that for each $p \in F$, C_p is a model of T. By definition of $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{pfc}}$, we know that in $\mathbb{A}_p$, we have $A \mathop{\smile\!\!\!\!\!\mid}_{C}^{p} B$, $a_0 \mathop{\smile\!\!\!\!\!\mid}_{C}^{p} A$ and $a_1 \mathop{\smile\!\!\!\!\!\mid}_{C}^{p} B$. As T is simple, the independence theorem over a model implies that there is some tuple a_p in $\mathbb{A}_p$ such that $a_p \equiv_{AC}^{L} a_0$, $a_p \equiv_{BC}^{L} a_1$ and $a_p \mathop{\smile\!\!\!\!\!\mid}_{C}^{p} AB$. For each $p \in F$, let $q^p(x) = \mathrm{tp}_L(a_p/ABC)$ considered as an L-type in $\mathbb{A}_p$. By Lemma 1.6.13, denoting the relativization of q^p to the parametrized language with respect to p by q^p_p, we know that the type $\bigcup_{p \in F} q^p_p$ is consistent. Let a be a realization. Then $a \equiv_{AC} a_0$ and $a \equiv_{BC} a_1$ in $\mathbb{A}_p$ for all $p \in F$ so $a \equiv_{ACF} a_0$ and $a \equiv_{BCF} a_1$.

Now we solve the problem for parameters. First assume that b_0, b_1 are *singletons* in $\mathbb{B}$. Without loss of generality $b_0, b_1 \notin F$ (as otherwise they are equal by assumption, and there is nothing to do). By quantifier-elimination, we need some $b \notin D \cup E \cup F$ so that $\langle aAC \rangle_b \cong \langle a_0AC \rangle_{b_0}$ and $\langle aBC \rangle_b \cong \langle a_1BC \rangle_{b_1}$. First, find $b_2 \equiv_{ACF} b_0$ and $b_3 \equiv_{BCF} b_1$ outside of $D \cup E \cup F$ so that $\langle aAC \rangle_{b_2} \cong \langle a_0AC \rangle_{b_0}$ and $\langle aBC \rangle_{b_3} \cong \langle a_1BC \rangle_{b_1}$. So $ab_2 \equiv_{ACF} a_0b_0$ and $ab_3 \equiv_{BCF} a_1b_1$. Now $b_2 \equiv_{aCF} b_3$ and $aAC \cap aBC \subseteq aC$, so Lemma 1.6.14 applies and we can find a b so that $\langle aAC \rangle_b \cong \langle aAC \rangle_{b_2}$ and $\langle aBC \rangle_b \cong \langle aBC \rangle_{b_3}$, and we can take this b to be outside of $D \cup E \cup F$. Now as $b \notin D \cup E \cup F$, we have $ab \equiv_{ACDF} a_0b_0$ and $ab \equiv_{BCEF} a_1b_1$.

Now let $b_0 = (b_{0,i} : i < k), b_1 = (b_{1,i} : i < k)$ be arbitrary tuples from $\mathbb{B}$. Without loss of generality, all of the elements in $\{b_{t,i} : i < k\}$ are pairwise-distinct, for $t \in \{0, 1\}$. Let $S_t = \{i < k : b_{t,i} \notin F\}$ for $t \in \{0, 1\}$, note that $S_0 = S_1 = S$ as $b_0 \equiv_F b_1$. Repeatedly applying the argument above for singletons, we can find *pairwise distinct* b'_i for $i \in S$ such that $a, b'_i \equiv_{ACDF} a_0, b_{0,i}$ and $a, b'_i \equiv_{BCEF} a_1, b_{1,i}$ for all $i \in S$. Let $b^* = (b^*_i : i < k)$ be defined by taking $b^*_i = b_{0,i} = b_{1,i}$ for all $i \notin S$ and $b^*_i = b'_i$ for all $i \in S$. As there are no relations in the language involving more than one element from the parameter sort except for the equality, it follows that $a, b^* \equiv_{ACDF} a_0, b_0$ and $a, b^* \equiv_{BCEF} a_1, b_1$ — as wanted.

□

Theorem 1.6.17. *Assume T is simple. Then $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{pfc}}$ is an Aut($\mathbb{M}$)-invariant independence relation on small subsets of the monster $\mathbb{M} \models T_{\mathrm{pfc}}$ such that it satisfies, for an arbitrary $M \models T_{\mathrm{pfc}}$:*

1. *strong finite character: if $a \not\mathop{\smile\!\!\!\!\!\mid}_{M}^{\mathrm{pfc}} b$, then there is a formula $\varphi(x, b, m) \in tp(a/bM)$ such that for any $a' \models \varphi(x, b, m)$, $a' \not\mathop{\smile\!\!\!\!\!\mid}_{M}^{\mathrm{pfc}} b$;*

2. *existence over models:* $M \models T_{\mathrm{pfc}}$ *implies* $a \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} M$ *for any* a;

3. *monotonicity:* $aa' \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} bb' \implies a \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} b$;

4. *symmetry:* $a \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} b \iff b \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} a$;

5. *independent amalgamation:* $c_0 \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} c_1$, $b_0 \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} c_0$, $b_1 \mathop{\smile\!\!\!\!\!\!|}_M^{\mathrm{pfc}} c_1$, $b_0 \equiv_M b_1$ *implies there exists* b *with* $b \equiv_{c_0 M} b_0$, $b \equiv_{c_1 M} b_1$.

Proof. Automorphism invariance and (1)-(4) are immediate from the definition of $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{pfc}}$, using that T is simple and hence non-forking independence satisfies all these properties; (5) was proven in Proposition 1.6.16. □

Corollary 1.6.18. *Suppose T is a simple theory which is the theory of a Fraïssé limit of a Fraïssé class $\mathbb{K}$ satisfying SAP. Then T_{pfc} is NSOP$_1$. Moreover, if the D-rank of T is ≥ 2, then T_{pfc} is not simple.*

Proof. By Proposition 6.2.4, T_{pfc} is NSOP$_1$, as $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{pfc}}$ gives an independence relation satisfying all the hypotheses. So now we prove that T_{pfc} is not simple, under the assumption that the D-rank of T is ≥ 2. This assumption implies that there is an L-formula $\varphi(x;y)$ and an indiscernible sequence $(a_i)_{i<\omega}$ so that $\{\varphi(x;a_i) : i < \omega\}$ is k-inconsistent for some k and the set defined by $\varphi(x;a_i)$ is infinite. Let $M \models T$ be some model containing the sequence $(a_i)_{i<\omega}$. Construct an L_{pfc}-structure N with domain $\omega \sqcup M$ and relations interpreted so that $N \models R_i(b) \iff M \models R(b)$ for each tuple $b \in M$, every $i < \omega$, and relation symbol R of L. Extend N to $\tilde{N} \models T_{\mathrm{pfc}}$. Let $\psi(x;y,z)$ be the formula $\varphi_z(x;y)$ and define an array $(b_{ij})_{i,j<\omega}$ by $b_{ij} = (a_j, i) \in M \times \omega \subset \tilde{N}^2$. Using Lemma 1.6.13, it is easy to check that for all $f : \omega \to \omega$, $\bigcup_{i<\omega}\{\psi(x;b_{if(i)})\}$ is consistent. Also $\{\psi(x;b_{ij}) : j < \omega\}$ is k-inconsistent for all i so ψ witnesses TP$_2$. □

Remark 1.6.19. For the above argument to work, we used that the formula witnessing dividing was non-algebraic — this fails in many natural examples (e.g. the random graph). However, given an L-structure M, define the *imaginary cover* of M as follows: let L' be the language L together with a new binary relation symbol E for an equivalence relation, and let $\tilde{M}$ be the L'-structure obtained by replacing each element of M with an infinite E-class and defining the relations of L on $\tilde{M}$ on the corresponding E-classes. Now it is easy to check that Age($\tilde{M}$) has SAP, the theory of $\tilde{M}$ is simple of D-rank at least 2.

Corollary 1.6.20. *T^*_{feq} is NSOP$_1$.*

Proof. The theory T of an equivalence relation with infinitely many infinite classes is a stable theory, obtained as the Fraïssé limit of all finite models of the theory of an equivalence relation. This class has no algebraicity, so it satisfies SAP. T_{pfc} is exactly T^*_{feq}, so it is NSOP$_1$. □

This result was claimed in [SU08], but the proof is apparently incorrect due to an illegitimate use of tree-indiscernibles. See the footnote on [HS14, p. 22] for a discussion.

Theories approximated by simple theories

In her thesis [HS14], Gwyneth Harrison-Shermoen considers theories that have a model approximated by a directed system $\mathcal{H}$ of homogeneous substructures, each of which has a simple theory. She proves that such theories carry an invariant independence notion $\mathop{\smile\hskip-0.9em\raise0.2ex\hbox{$\vert$}}^{\lim}$ satisfying strong finite character, monotonicity, symmetry, and existence over a model (existence over a model is implied by Claim 3.3.4 in [HS14]). Finally, she observes that if non-forking independence $\mathop{\smile\hskip-0.9em\raise0.2ex\hbox{$\vert$}}^{f}$ satisfies the independence theorem over algebraically closed sets for each model in $\mathcal{H}$, then so does $\mathop{\smile\hskip-0.9em\raise0.2ex\hbox{$\vert$}}^{lim}$ for the approximated theory. Hence, we obtain the following:

Corollary 1.6.21. *Suppose T is a theory approximated, in the sense of Harrison-Shermoen, by a directed system of structures each with a simple theory in which $\mathop{\smile\hskip-0.9em\raise0.2ex\hbox{$\vert$}}^{f}$ satisfies the independence theorem over algebraically closed sets. Then T is NSOP$_1$.*

1.7 Lemmas on preservation of indiscernibility

Lemma 1.7.1. *Suppose $\eta_0, \ldots, \eta_{l-1}$, $\nu_0, \ldots, \nu_{l-1}$ are elements of $\omega^{<\omega}$. Let $\overline{\eta}$ and $\overline{\nu}$ denote enumerations of the $\wedge$-closures of $\{\eta_i : i < l\}$ and $\{\nu_i : i < l\}$ respectively. Then if*

$$qftp_{L_s}(\eta_0, \ldots, \eta_{l-1}) = qftp_{L_s}(\nu_0, \ldots, \nu_{l-1}),$$

then

$$qftp_{L_s}(\overline{\eta}) = qftp_{L_s}(\overline{\nu}).$$

Proof. Easy. See Remark 3.2 from [KKS14] □

Lemma 1.7.2. *Let $\eta_0, \ldots, \eta_{l-1}, \nu_0, \ldots, \nu_{l-1} \in \omega^{<\omega}$ be such that*

$$qftp_{L_s}(\eta_0, \ldots, \eta_{l-1}) = qftp_{L_s}(\nu_0, \ldots, \nu_{l-1}).$$

Suppose $i < l$ and $\eta \lhd \eta_i$, $\nu \lhd \nu_i$ with $l(\eta) = l(\nu)$. Then, setting $\eta_l = \eta$ and $\nu_l = \nu$, we have

$$qftp_{L_s}(\eta_0, \ldots, \eta_l) = qftp_{L_s}(\nu_0, \ldots, \nu_l).$$

Proof. Without loss of generality, we may take $\{\eta_i : i < l\}$ and $\{\nu_i : i < l\}$ to be $\wedge$-closed, by the previous lemma. Then $\{\eta_i : i < l+1\}$ and $\{\nu_i : i < l+1\}$ are also $\wedge$-closed. So we need only to check that for any $j, j' < l+1$

1. $\eta_j \lhd \eta_{j'} \iff \nu_j \iff \nu_j'$
2. $\eta_j <_{lex} \eta_{j'} \iff \nu_j <_{lex} \nu_{j'}$

We have 3 cases.

Case 1: $j, j' < l$.

(1) and (2) follow by assumption.

Case 2: $j < l$ and $j' = l$

$$\begin{aligned}
\eta_j \lhd \eta_l &\iff \eta_j \lhd \eta_i \text{ and } l(\eta_j) \leq l(\eta_l) \\
&\iff \eta_j \lhd \eta_i \wedge \bigvee_{k<l(\eta_l)} P_k(\eta_j) \\
&\iff \nu_j \lhd \nu_i \wedge \bigvee_{k<l(\nu_l)} P_k(\nu_j) \\
&\iff \nu_j \lhd \nu_l. \\
\eta_j <_{lex} \eta_l &\iff l(\eta_j \wedge \eta_i) < l(\eta_l) \text{ and } \eta_j <_{lex} \eta_i \\
&\iff \left(\bigvee_{k<l(\eta_l)} P_k(\eta_j \wedge \eta_i) \right) \wedge \eta_j <_{lex} \eta_i \\
&\iff \left(\bigvee_{k<l(\nu_l)} P_k(\nu_j \wedge \nu_i) \right) \wedge \nu_j <_{lex} \nu_i \\
&\iff \nu_l <_{lex} \nu_j.
\end{aligned}$$

Case 3: $j = l$ and $j' < l$

$$\begin{aligned}
\eta_l \lhd \eta_j &\iff \eta_l \lhd (\eta_i \wedge \eta_j) \\
&\iff \bigvee_{l(\eta_l)<k\leq l(\eta_i)} P_k((\eta_i \wedge \eta_j)) \\
&\iff \bigvee_{l(\nu_l)<k\leq l(\nu_i)} P_k((\nu_i \wedge \nu_j)) \\
&\iff \nu_l \lhd \nu_j \\
\eta_l <_{lex} \eta_j &\iff (l(\eta_j \wedge \eta_i) < l(\eta_l)) \to \eta_i <_{lex} \eta_j \\
&\iff \left(\bigvee_{k<l(\eta_l)} P_k(\eta_j \wedge \eta_i) \right) \to \eta_i <_{lex} \eta_j \\
&\iff \left(\bigvee_{k<l(\nu_l)} P_k(\nu_j \wedge \nu_i) \right) \to \nu_i <_{lex} \nu_j \\
&\iff \nu_l <_{lex} \nu_j.
\end{aligned}$$

□

Lemma 1.7.3. *Let $(a_\eta)_{\eta\in\omega^{<\omega}}$ be an s-indiscernible tree. If $(a'_\eta)_{\eta\in\omega^{<\omega}}$ is the k-fold widening of $(a_\eta)_{\eta\in\omega^{<\omega}}$ at level n, then $(a'_\eta)_{\eta\in\omega^{<\omega}}$ is also s-indiscernible.*

Proof. Pick $\eta_0, \ldots, \eta_{l-1}$ and $\nu_0, \ldots, \nu_{l-1}$ in $\omega^{<\omega}$ so that

$$\mathrm{qftp}_{L_s}(\eta_0, \ldots, \eta_{l-1}) = \mathrm{qftp}_{L_s}(\nu_0, \ldots, \nu_{l-1}).$$

By Lemma 1.7.2, we may assume that $\{\eta_i : i < l\}$ and $\{\nu_i : i < l\}$ are both $\wedge$-closed and closed under initial segment. Moreover, we may assume that these elements have been enumerated so that for some $m \le l$, $l(\eta_i), l(\nu_i) < n$ if and only if $i \ge m$. So for each $i < m$, we may write

$$\begin{aligned} \eta_i &= \mu_i \frown \alpha_i \frown \xi_i \\ \nu_i &= \upsilon_i \frown \beta_i \frown \rho_i, \end{aligned}$$

where $\mu_i, \upsilon_i \in \omega^{n-1}$, $\alpha_i, \beta_i \in \omega$, and $\xi_i, \rho_i \in \omega^{<\omega}$. For each $i < m$, let

$$\begin{aligned} \overline{\eta}_i &= (\mu_i \frown (k\alpha_i) \frown \xi_i, \mu_i \frown (k\alpha_i + 1) \frown \xi_i, \ldots, \mu_i \frown (k\alpha_i + k - 1) \frown \xi_i) \\ \overline{\nu}_i &= (\upsilon_i \frown (k\beta_i) \frown \rho_i, \upsilon_i \frown (k\beta_i + 1) \frown \rho_i, \ldots, \upsilon_i \frown (k\beta_i + k - 1) \frown \rho_i). \end{aligned}$$

and for $m \le i < l$, let $\overline{\eta}_i = \eta_i$, $\overline{\nu}_i = \nu_i$. Now we must show that

$$\mathrm{qftp}_{L_s}(\overline{\eta}_0, \ldots, \overline{\eta}_{l-1}) = \mathrm{qftp}_{L_s}(\overline{\nu}_0, \ldots, \overline{\nu}_{l-1}).$$

It is clear that the sets $\bigcup_{i<l} \overline{\eta}_i$ and $\bigcup_{i<l} \overline{\nu}_i$ are closed under initial segment. They are also closed under $\wedge$: this is obvious for elements of length $< n$ and for elements of longer length whose meet has length $< n$ by our assumptions. On the other hand if, for some $i, i' < l$ and $j, j' < k$, $l((\overline{\eta}_i)_j), l((\overline{\nu}_{i'})_{j'}) \ge n$ and $l((\overline{\eta}_i)_j \wedge (\overline{\nu}_{i'})_{j'}) \ge n$, then if $j = j'$, we have $(\overline{\eta}_i)_j \wedge (\overline{\nu}_{i'})_{j'} = \overline{(\eta_i \wedge \eta_{i'})}_j$ and if $j \ne j'$, then $(\overline{\eta}_i)_j \wedge (\overline{\nu}_{i'})_{j'}$ is equal to the common initial segment of each element of length $n-1$. In the first case, the meet is enumerated in one of the tuples because our initial set of tuples was $\wedge$-closed, in the second case because it was taken to be closed under initial segment. To check equality of the quantifier-free types, we have 3 cases:

Case 1: $i, i' \ge m$ Follows by assumption, as for any $i \ge m$, $\overline{\eta}_i = \eta_i$ and $\overline{\nu}_i = \nu_i$.

Case 2: $i \ge m$, $i' < m$ and $j < k$

$$\begin{aligned} \overline{\eta}_i \lhd (\overline{\eta}_{i'})_j &\iff \overline{\nu}_i \lhd (\overline{\nu}_{i'})_j \\ \overline{\eta}_i <_{lex} (\overline{\eta}_{i'})_j &\iff \overline{\nu}_i <_{lex} (\overline{\eta}_{i'})_j \\ (\overline{\eta}_{i'})_j <_{lex} \overline{\eta}_i &\iff (\overline{\nu}_{i'})_j <_{lex} \overline{\nu}_i \end{aligned}$$

Case 3: $i, i' < m$ and $j, j' < k$

$$\begin{aligned}
(\overline{\eta}_i)_j \lhd (\overline{\eta}_{i'})_{j'} &\iff \eta_i \lhd \eta_{i'} \text{ and } j = j' \\
&\iff \nu_i \lhd \nu_{i'} \text{ and } j = j' \\
&\iff (\overline{\nu}_i)_j \lhd (\overline{\nu}_{i'})_{j'} \\
(\overline{\eta}_i)_j <_{lex} (\overline{\eta}_{i'})_{j'} &\iff (\eta_i <_{lex} \eta_j \text{ and } (l(\eta_i \wedge \eta_j) < n \text{ or } j = j')) \text{ or} \\
&\quad\;\; (l(\eta_i \wedge \eta_{i'}) \geq n \text{ and } j < j') \\
&\iff (\nu_i <_{lex} \nu_j \text{ and } (l(\nu_i \wedge \nu_j) < n \text{ or } j = j')) \text{ or} \\
&\quad\;\; (l(\nu_i \wedge \nu_{i'}) \geq n \text{ and } j < j') \\
&\iff (\overline{\nu}_i)_j <_{lex} (\overline{\nu}_{i'})_{j'}.
\end{aligned}$$

□

Lemma 1.7.4. *Let $(a_\eta)_{\eta\in\omega^{<\omega}}$ be an s-indiscernible tree. If $(a''_\eta)_{\eta\in\omega^{<\omega}}$ is the k-fold stretch of $(a_\eta)_{\eta\in\omega^{<\omega}}$ at level n, then $(a''_\eta)_{\eta\in\omega^{<\omega}}$ is also s-indiscernible.*

Proof. Given $\eta \in \omega^{<\omega}$, let

$$\overline{\eta} = \begin{cases} \eta & \text{if } l(\eta) < n \\ (\eta, \eta \frown 0, \ldots, \eta \frown 0^{k-1}) & \text{if } l(\eta) = n \\ \nu \frown 0^{k-1} \frown \xi & \text{if } \eta = \nu \frown \xi, \text{ with } \nu \in \omega^n, \xi \neq \emptyset \end{cases}$$

Pick $\eta_0, \ldots, \eta_{l-1}, \nu_0, \ldots, \nu_{l-1} \in \omega^{<\omega}$ so that

$$\mathrm{qftp}_{L_s}(\eta_0, \ldots, \eta_{l-1}) = \mathrm{qftp}_{L_s}(\nu_0, \ldots, \nu_{l-1}),$$

and, without loss of generality, we may suppose $\{\eta_i : i < l\}$ and $\{\nu_i : i < l\}$ are both $\wedge$-closed. We must show that

$$\mathrm{qftp}_{L_s}(\overline{\eta}_0, \ldots, \overline{\eta}_{l-1}) = \mathrm{qftp}_{L_s}(\overline{\nu}_0, \ldots, \overline{\nu}_{l-1}).$$

Assume that $\{\overline{\eta}_i : i < l\}$ is ordered so that $i < m$ if and only if $l(\eta_i) = n$, and similarly for $\{\overline{\nu}_i : i < l\}$. Clearly $\{\overline{\eta}_i : i < l\}$ and $\{\overline{\nu}_i : i < l\}$ are also $\wedge$-closed, so we have to check that the two sequences of tuples have the same quantifier type with respect to the relations $<_{lex}$ and $\lhd$. We'll show this by considering the various cases:

Case 1: $i, i' \geq m$. Then

$$\begin{aligned}
\overline{\eta}_i \lhd \overline{\eta}_{i'} &\iff \eta_i \lhd \eta_{i'} \\
&\iff \nu_i \lhd \nu_{i'} \\
&\iff \overline{\nu}_i \lhd \overline{\nu}_{i'} \\
\overline{\eta}_i <_{lex} \overline{\eta}_i &\iff \eta_i <_{lex} \eta_{i'} \\
&\iff \nu_i <_{lex} \nu_{i'} \\
&\iff \overline{\nu}_i <_{lex} \overline{\nu}_{i'}.
\end{aligned}$$

Case 2: $i, i' < m$ and $j, j' < k$. Then

$$\begin{aligned}
(\overline{\eta}_i)_j \lhd (\overline{\eta}_{i'})_{j'} &\iff (\eta_i = \eta_{i'}) \wedge (j < j') \\
&\iff (\nu_i = \nu_{i'}) \wedge (j < j') \\
&\iff (\overline{\nu}_i)_j \lhd (\overline{\nu})_{j'} \\
(\overline{\eta}_i)_j <_{lex} (\overline{\eta}_{i'})_{j'} &\iff \eta_i <_{lex} \eta_{i'} \vee (\nu_i = \nu_{i'} \wedge j < j') \\
&\iff \nu_i <_{lex} \nu_{i'} \vee (\nu_i = \nu_{i'} \wedge j < j') \\
&\iff (\overline{\nu}_i)_j <_{lex} (\overline{\nu}_{i'})_{j'}.
\end{aligned}$$

Case 3: $i < m$, $i' \geq m$, $j < k$.

$$\begin{aligned}
(\overline{\eta}_i)_j \lhd \overline{\eta}_{i'} &\iff \eta_i \lhd \eta_{i'} \\
&\iff \nu_i \lhd \nu_{i'} \\
&\iff (\overline{\nu}_i)_j \lhd \overline{\nu}_i \\
\overline{\eta}_{i'} \lhd (\overline{\eta}_i)_j &\iff \eta_{i'} \lhd \eta_i \\
&\iff \nu_{i'} \lhd \nu_i \\
&\iff (\overline{\nu}_{i'})_j \lhd \overline{\nu}_i \\
(\overline{\eta}_i)_j <_{lex} \overline{\eta}_{i'} &\iff \eta_i <_{lex} \eta_{i'} \\
&\iff \nu_i <_{lex} \nu_{i'} \\
&\iff (\overline{\nu}_i)_j <_{lex} \overline{\nu}_{i'} \\
\overline{\eta}_{i'} <_{lex} (\overline{\eta}_i)_j &\iff \nu_{i'} <_{lex} \nu_i \\
&\iff \nu_{i'} <_{lex} \nu_i \\
&\iff \overline{\nu}_{i'} <_{lex} (\overline{\nu}_i)_j.
\end{aligned}$$

□

Lemma 1.7.5. *1. Each tuple $a_\eta^{(n)}$ may be enumerated as $(a_{\nu\frown\eta} : \nu \in 2^n)$*

2. If $(a_\eta)_{\eta\in 2^{<\kappa}}$ is strongly indiscernible, then for all n, the n-fold fattening $(a_\eta^{(n)})_{\eta\in 2^{<\kappa}}$ is strongly indiscernible over C_n

Proof. (1) This is trivial for $n = 0$. Then if true for n, we have

$$a_\eta^{(n+1)} = (a_{0\frown\eta}^{(n)}, a_{1\frown\eta}^{(n)}) = ((a_{\nu\frown 0\frown\eta} : \nu \in 2^n), (a_{\nu\frown 1\frown\eta} : \nu \in 2^n)) = (a_{\xi\frown\eta} : \xi \in 2^{n+1}).$$

(2) By (1) we have $a_\eta^{(n+1)} = (a_{\mu\frown\eta} : \mu \in 2^n)$. Let $\overline{\mu} = (\mu \in 2^{\leq n})$. In order to show indiscernibility over C_n have to show that if $\eta_0, \dots, \eta_{k-1}, \nu_0, \dots, \nu_{k-1} \in 2^{<\omega}$ and

$$\mathrm{qftp}_{L_0}(\eta_0, \dots, \eta_{k-1}) = \mathrm{qftp}_{L_0}(\nu_0, \dots, \nu_{k-1})$$

then $\mathrm{qftp}_{L_0}(\overline{\mu}, (a_{\mu\frown\eta_0} : \mu \in 2^n), \dots, (a_{\mu\frown\eta_{k-1}} : \mu \in 2^n))$ is equal to $\mathrm{qftp}_{L_0}(\overline{\mu}, (a_{\mu\frown\nu_0} : \mu \in 2^n), \dots, (a_{\mu\frown\nu_{k-1}} : \mu \in 2^n))$. To this end, we may assume $\{\eta_0, \dots, \eta_{k-1}\}$ and $\{\nu_0, \dots, \nu_{k-1}\}$

are meet-closed. Then $2^{\leq n} \cup \{\mu \frown \eta_i : \mu \in 2^n, i < k\}$ and $2^{\leq n} \cup \{\mu \frown \nu_i : \mu \in 2^n, i < k\}$ is also meet-closed and we just have to check that the tuples in the above equation have the same time with respect to the language $L_t = \{\lhd, <_{lex}\}$. Choose ξ_0, ξ_1 from the tuple $(\overline{\mu}, (a_{\mu \frown \eta_0} : \mu \in 2^n), \ldots, (a_{\mu \frown \eta_{k-1}} : \mu \in 2^n))$ and ρ_0, ρ_1 from $(\overline{\mu}, (a_{\mu \frown \nu_0} : \mu \in 2^n), \ldots, (a_{\mu \frown \nu_{k-1}} : \mu \in 2^n))$ so that ξ_i sits in the same position in the enumeration of the tuple as ρ_i for $i = 0, 1$. Now, we must show that $\xi_0 <_{lex} \xi_1$ if and only if $\rho_0 <_{lex} \rho_1$ and $\xi_0 \unlhd \xi_1$ if and only if $\rho_0 \unlhd \rho_1$. Choose arbitrary $\mu_0, \mu_1 \in 2^{\leq n}$, η_i, η_j, ν_i, ν_j.

Case 1: $l(\mu_0) = l(\mu_1) = n$, $\xi_0 = \mu_0 \frown \eta_i$, $\xi_1 = \mu_1 \frown \eta_j$, and hence $\rho_0 = \mu_0 \frown \nu_i$ and $\rho_1 = \mu_1 \frown \nu_j$.

$$
\begin{aligned}
\mu_0 \frown \eta_i \unlhd \mu_1 \frown \eta_j \quad &\iff \quad \mu_0 = \mu_1 \wedge \eta_i \unlhd \eta_j \\
&\iff \quad \mu_0 = \mu_1 \wedge \nu_i \unlhd \nu_j \\
&\iff \quad \mu_0 \frown \nu_i \lhd \mu_1 \frown \nu_j \\
\mu_0 \frown \eta_i <_{lex} \mu_1 \frown \eta_j \quad &\iff \quad \mu_0 <_{lex} \mu_1 \vee (\mu_0 = \mu_1 \wedge \eta_i <_{lex} \eta_j) \\
&\iff \quad \mu_0 <_{lex} \mu_1 \vee (\mu_0 = \mu_1 \wedge \nu_i <_{lex} \nu_{i'}) \\
&\iff \quad \mu_0 \frown \nu_i <_{lex} \mu_1 \frown \nu_j
\end{aligned}
$$

Case 2: $\xi_0 = \mu_0$, $\xi_1 = \mu_1$, $\rho_0 = \mu_0$, and $\rho_1 = \mu_1$.
Clear.
Case 3: $l(\mu_0) = n$, $\xi_0 = \mu_0 \frown \eta_i$, $\xi_1 = \mu_1$, $\rho_0 = \mu_0 \frown \nu_i$, $\rho_1 = \mu_1$.
It is never the case that $\mu_0 \frown \eta_i \lhd \mu_1$ or $\mu_0 \frown \nu_i \lhd \mu_1$ so it suffices to check $<_{lex}$:

$$
\begin{aligned}
\mu_0 \frown \eta_i <_{lex} \mu_1 \quad &\iff \quad \mu_0 <_{lex} \mu_1 \\
&\iff \quad \mu_0 \frown \nu_i <_{lex} \mu_1.
\end{aligned}
$$

Case 4: $l(\mu_1) = n$, $\xi_0 = \mu_0$, $\xi_1 = \mu_1 \frown \nu_j$, $\rho_0 = \mu_0$, $\rho_1 = \mu_1 \frown \nu_j$.

$$
\begin{aligned}
\mu_0 \unlhd \mu_1 \frown \eta_j \quad &\iff \quad \mu_0 \unlhd \mu_1 \\
&\iff \quad \mu_0 \unlhd \mu_1 \frown \nu_j \\
\mu_0 \leq_{lex} \mu_1 \frown \eta_j \quad &\iff \quad \mu_0 \leq_{lex} \mu_1 \\
&\iff \quad \mu_0 \leq_{lex} \mu_1 \frown \nu_j
\end{aligned}
$$

□

Lemma 1.7.6. *If $(a_\eta)_{\eta \in 2^{<\omega}}$ is strongly indiscernible, then for all natural numbers $k \geq 1$, the k-fold elongation $(a'_\eta)_{\eta \in 2^{<\omega}}$ of $(a_\eta)_{\eta \in 2^{<\omega}}$ is also strongly indiscernible.*

Proof. Given $\eta \in 2^{<\omega}$, with $l(\eta) = n$, we defined $\tilde{\eta} \in 2^{<\omega}$ to be the element with length $k(l(\eta) - 1) + 1$ defined by

$$
\tilde{\eta}(i) = \begin{cases} \eta(i/k) & \text{if } k|i \\ 0 & \text{otherwise} \end{cases}
$$

As the *k-fold elongation* of $(a_\eta)_{\eta\in 2^{<\omega}}$ is defined to be the tree $(b_\eta)_{\eta\in 2^{<\omega}}$ where

$$b_\eta = (a_{\tilde{\eta}}, a_{\tilde{\eta}\frown 0}, \ldots, a_{\tilde{\eta}\frown 0^{k-1}}).$$

Write $\overline{\eta}$ for the tuple $(\tilde{\eta}, \tilde{\eta} \frown 0, \ldots, \tilde{\eta} \frown 0^{k-1})$. We are reduced to showing that if $\eta_0, \ldots, \eta_{l-1}$, $\nu_0, \ldots, \nu_{l-1}$ are elements of $2^{<\omega}$ so that

$$\mathrm{qftp}_{L_0}(\eta_0, \ldots, \eta_{l-1}) = \mathrm{qftp}_{L_0}(\nu_0, \ldots, \nu_{l-1})$$

then

$$\mathrm{qftp}_{L_0}(\overline{\eta}_0, \ldots, \overline{\eta}_{l-1}) = \mathrm{qftp}_{L_0}(\overline{\nu}_0, \ldots, \overline{\nu}_{l-1}).$$

We may assume that $\{\eta_i : i < l\}$ and $\{\nu_i : i < l\}$ are both $\wedge$-closed, from which it follows that $\{\overline{\eta}_i : i < l\}$ and $\{\overline{\nu}_i : i < l\}$ are both $\wedge$-closed. So we must check that $(\overline{\eta}_i : i < l)$ and $(\overline{\nu}_i : i < l)$ have the same quantifier-free type with respect to the language $L_t = \langle \trianglelefteq, <_{lex} \rangle$. We note

$$\begin{aligned}
\tilde{\eta}_i \frown 0^l \trianglelefteq \tilde{\eta}_j \frown 0^{l'} &\iff \tilde{\eta}_i \triangleleft \tilde{\eta}_j \vee (\tilde{\eta}_i = \tilde{\eta}_j \wedge l \leq l') \\
&\iff \eta_i \triangleleft \eta_j \vee (\eta_i = \eta_j \wedge l \leq l') \\
&\iff \nu_i \triangleleft \nu_j \vee (\nu_i = \nu_j \wedge l \leq l') \\
&\iff \tilde{\nu}_i \triangleleft \tilde{\nu}_j \vee (\tilde{\nu}_i = \tilde{\nu}_j \wedge l \leq l') \\
&\iff \tilde{\nu}_i \frown 0^l \triangleleft \tilde{\nu}_j \frown 0^{l'} \\
\tilde{\eta}_i \frown 0^l <_{lex} \tilde{\eta}_j \frown 0^{l'} &\iff \tilde{\eta}_i <_{lex} \tilde{\eta}_j \vee (\tilde{\eta}_i = \tilde{\eta}_j \wedge l < l') \\
&\iff \eta_i <_{lex} \eta_j \vee (\eta_i = \eta_j \wedge l < l') \\
&\iff \nu_i <_{lex} \nu_j \vee (\nu_i = \nu_j \wedge l < l') \\
&\iff \tilde{\nu}_i <_{lex} \tilde{\nu}_j \vee (\tilde{\nu}_i = \tilde{\nu}_j \wedge l < l') \\
&\iff \tilde{\nu}_i \frown 0^l <_{lex} \tilde{\nu}_j \frown 0^{l'}.
\end{aligned}$$

□

Lemma 1.7.7. *Suppose $(a_\eta)_{\eta\in 2^{<\omega}}$ is a strongly indiscernible tree over C.*

1. *Define a function $h : 2^{<\omega} \to 2^{<\omega}$ by $h(\emptyset) = \emptyset$ and $h(\eta) = h(\nu) \frown 0 \frown \langle i \rangle$ whenever $\eta = \nu \frown \langle i \rangle$. Then $(a_{h(\eta)})_{\eta\in 2^{<\omega}}$ is strongly indiscernible over C.*

2. *For each n, define a map $h_n : 2^{<\omega} \to 2^{<\omega}$ by*

$$h_n(\eta) = \begin{cases} h(\eta) & \text{if } l(\eta) \leq n \\ h(\nu) \frown \xi & \text{if } \eta = \nu \frown \xi, l(\nu) = n. \end{cases}$$

Then $(a_{h_n(\eta)})_{\eta\in 2^{<\omega}}$ is strongly indiscernible over C.

Proof. (1) At the outset, we note that $\eta \trianglelefteq \nu \iff h(\eta) \trianglelefteq h(\nu)$ and $\eta <_{lex} \nu \iff h(\eta) <_{lex} h(\nu)$. The only difficulty arises from $\wedge$ which is not preserved by h, because if $\eta \perp \nu$ and $\eta \wedge \nu = \xi$ then $h(\eta) \wedge h(\nu) = h(\xi) \frown 0$.

It suffices to show that if $\overline{\eta}, \overline{\nu}$ are finite tuples from $2^{<\omega}$ with $\mathrm{qftp}_{L_0}(\overline{\eta}) = \mathrm{qftp}_{L_0}(\overline{\nu})$ then $\mathrm{qftp}_{L_0}(h(\overline{\eta})) = \mathrm{qftp}_{L_0}(h(\overline{\nu}))$. Given such $\overline{\eta}, \overline{\nu}$, it is clear that if $\mathrm{qftp}_{L_0}(h(\overline{\eta})) \neq \mathrm{qftp}_{L_0}(h(\overline{\nu}))$ then $\mathrm{qftp}_{L_0}(h(\overline{\eta}')) \neq \mathrm{qftp}_{L_0}(h(\overline{\nu}'))$ where $\overline{\eta}'$ and $\overline{\nu}'$ are the $\wedge$-closures of $\overline{\eta}$ and $\overline{\nu}$ respectively. So we may assume $\overline{\eta}$ and $\overline{\nu}$ are $\wedge$-closed. We may assume that the tuple $\overline{\eta} = \langle \eta_i : i < k \rangle$ is enumerated so that for some $l \leq k$, if $i < l$, then there are $\eta_j \perp \eta_{j'}$ so that $\eta_j \wedge \eta_{j'} = \eta_i$. It follows that the $\wedge$-closure of $h(\overline{\eta})$ may be enumerated as $\langle h(\eta_i) : i < k \rangle \frown \langle h(\eta_i) \frown 0 : i < l \rangle$, and, likewise, the $\wedge$-closure of $h(\overline{\nu})$ can be enumerated as $\langle h(\nu_i) : i < k \rangle \frown \langle h(\nu_i) \frown 0 : i < l \rangle$. Now we note that, by definition of h, if $i, j < k$

$$
\begin{aligned}
h(\eta_i) \vartriangleleft h(\eta_j) \frown 0 &\iff h(\eta_i) \frown 0 \vartriangleleft h(\eta_j) \\
&\iff h(\eta_i) \frown 0 \vartriangleleft h(\eta_j) \frown 0 \\
&\iff h(\eta_i) \vartriangleleft h(\eta_j) \\
h(\eta_i) <_{lex} h(\eta_j) \frown 0 &\iff h(\eta_i) \frown 0 <_{lex} h(\eta_j) \\
&\iff h(\eta_i) \frown 0 <_{lex} h(\eta_j) \frown 0 \\
&\iff h(\eta_i) <_{lex} h(\eta_j)
\end{aligned}
$$

And similarly for ν_i, ν_j. As h respects $\vartriangleleft$ and $<_{lex}$, and $\mathrm{qftp}_{L_0}(\overline{\eta}) = \mathrm{qftp}_{L_0}(\overline{\nu})$, it follows that $\mathrm{qftp}_{L_0}(h(\overline{\eta})) = \mathrm{qftp}_{L_0}(h(\overline{\nu}))$.

(2) is entirely similar. $\square$

Chapter 2

Invariants related to the tree property

2.1 Introduction

One of the fundamental discoveries in stability theory is that stability is local – a theory is stable if and only if no formula has the order property. Among the stable theories, one can obtain a measure of complexity by associating to each theory T its *stability spectrum*, namely, the class of cardinals λ such that T is stable in λ. A classification of stability spectra was given by Shelah in [She90, Chapter 3]. Part of this analysis amounts showing that stable theories do not have the tree property and, consequently, that forking satisfies local character. But a crucial component of that work was studying the approximations to the tree property can exist in stable theories and what structural consequences they have. These approximations were measured by a cardinal invariant of the theory called $\kappa(T)$, and Shelah's stability spectrum theorem gives an explicit description of the cardinals in which a given theory T was stable in terms of the cardinality of the set of types over the empty set and $\kappa(T)$. Shelah used the definition of $\kappa(T)$ as a template for quantifying the global approximations to other tree properties in introducing the invariants $\kappa_{\text{cdt}}(T)$, $\kappa_{\text{sct}}(T)$, and $\kappa_{\text{inp}}(T)$ which bound approximations to the tree property (TP), the tree property of the first kind (TP_1), and the tree property of the second kind (TP_2), respectively. Eventually, the local condition that a theory does not have the tree property (*simplicity*), and the global condition that $\kappa(T) = \kappa_{cdt}(T) = \aleph_0$ (*supersimplicity*) proved to mark substantial dividing lines. These invariants provide a coarse measure of the complexity of the theory, providing a "quantitative" description of the patterns that can arise among forking formulas. They are likely to continue to play a role in the development of a structure theory for tame classes of non-simple theories.

Motivated by some questions from [She90], we explore which relationships that obtain between the *local* properties TP, TP_1, and TP_2 also hold for the *global* invariants $\kappa_{\text{cdt}}(T)$, $\kappa_{\text{sct}}(T)$, and $\kappa_{\text{inp}}(T)$. In short, we are pursuing the following analogy:

local	TP	TP_1	TP_2
global	κ_{cdt}	κ_{sct}	κ_{inp}

This continues the work done in Chapter 1, where, with Artem Chernikov, we considered a global analogue of the following theorem of Shelah:

Theorem. [She90, III.7.11] For complete theory T, $\kappa_{\mathrm{cdt}}(T) = \infty$ and only if $\kappa_{\mathrm{sct}}(T) = \infty$ or $\kappa_{\mathrm{inp}}(T) = \infty$. That is, T has the tree property if and only if it has the tree property of the first kind or the tree property of the second kind.

Shelah then asked if $\kappa_{\mathrm{cdt}}(T) = \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$ in general [She90, Question III.7.14][1]. In Chapter 1, we showed that is true under the assumption that T is countable. For a countable theory T, the only possible values of these invariants are $\aleph_0, \aleph_1$, and ∞ – our proof handled each cardinal separately using a different argument in each case. Here we consider this question without any hypothesis on the cardinality of T, answering the general question negatively:

Theorem. There is a theory T so that $\kappa_{\mathrm{cdt}}(T) \neq \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$. Moreover, it is consistent with ZFC that for every regular uncountable κ, there is a theory T with $|T| = \kappa$ and $\kappa_{\mathrm{cdt}}(T) \neq \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$.

To construct a theory T so that $\kappa_{\mathrm{cdt}}(T) \neq \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$, we use results on *strong colorings* constructed by Galvin under GCH and later by Shelah in ZFC. These results show that, at suitable regular cardinals, Ramsey's theorem fails in a particularly dramatic way. The statement $\kappa_{\mathrm{cdt}}(T) = \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$ amounts to saying that a certain large global configuration gives rise to another large configuration which is moreover very uniform. This has the feel of many statements in the partition calculus and we show that, in fact, a coloring $f : [\kappa]^2 \to 2$ can be used to construct a theory $T^*_{\kappa,f}$ such that the existence of a large inp- or sct-patterns relative to $T^*_{\kappa,f}$ implies some homogeneity for the coloring f. The theories built from the strong colorings of Galvin and Shelah, then, furnish ZFC counter-examples to Shelah's question, and also give a consistency result showing that, consistently, for every regular uncountable cardinal κ, there is a theory T with $|T| = \kappa$ and $\kappa_{\mathrm{cdt}}(T) \neq \kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T)$. This suggests that the aforementioned result of Chapter 1 for countable theories is in some sense the optimal result possible in ZFC.

Our second theorem is motivated by the following theorem of Shelah:

Theorem. [She90, VI.4.7] If T is not simple, $\mathcal{D}$ is a regular ultrafilter over I, M is an $|I|^{++}$-saturated model of T, then $M^I/\mathcal{D}$ is not $|I|^{++}$-compact.

[1]This formulation is somewhat inaccurate. Shelah defines for $x \in \{\mathrm{cdt}, \mathrm{inp}, \mathrm{sct}\}$, the cardinal invariant κr_x, which is the least regular cardinal $\geq \kappa_x$. Shelah's precise question was about the possible equality $\kappa r_{\mathrm{cdt}} = \kappa r_{\mathrm{sct}} + \kappa r_{\mathrm{inp}}$. For our purposes, we will only need to consider theories in which κ_x is a successor cardinal, so we will not need to distinguish between these two variations.

In an exercise, Shelah claims that the hypothesis that T is not simple in the above theorem may be replaced by the condition $\kappa_{\text{inp}}(T) > |I|^+$ and asks if $\kappa_{\text{cdt}}(T) > |I|^+$ suffices [She90, Question VI.4.20]. We prove the following:

Theorem. There is a theory T so $\kappa_{\text{inp}}(T) = \lambda^{++}$ yet for any regular ultrafilter $\mathcal{D}$ on λ and λ^{++}-saturated model of T, $M^\lambda/\mathcal{D}$ is λ^{++}-saturated. On the other hand, for an arbitrary T, if $\lambda = \lambda^{<\lambda}$ and $\kappa_{\text{sct}}(T) > \lambda^+$, M is an λ^{++}-saturated model of T and $\mathcal{D}$ is a regular ultrafilter over λ, then $M^\lambda/\mathcal{D}$ is not λ^{++}-compact.

This contradicts Shelah's Exercise VI.4.19 and *a fortiori* answers Question VI.4.20 negatively. Although $\kappa_{\text{inp}}(T) > |I|^+$ and hence $\kappa_{\text{cdt}}(T) > |I|^+$ do not suffice to guarantee a loss of saturation in the ultrapower, one can ask if $\kappa_{\text{sct}}(T) > |I|^+$ does suffice. Shelah's original argument for Theorem 2.5.4 does not generalize, but fortunately a recent new proof due to Malliaris and Shelah [MS15b] does and we point out how the revised question can be answered by an easy and direct adaptation of their argument. These results suggest that the rough-scale asymptotic structure revealed by studying the λ^{++}-compactness of ultrapowers on λ is global in nature and differs from the picture suggested by the local case considered by Shelah.

In order to construct these examples, it is necessary to build a theory capable of coding a complicated strong coloring yet simple enough that the invariants are still computable. This was accomplished by a method inspired by Medvedev's $\mathbb{Q}$ACFA construction [Med], realizing the theory as a union of theories in a system of finite reducts each of which is the theory of a Fraïssé limit. The theories in the finite reducts are $\aleph_0$-categorical and eliminate quantifiers and one may apply the Δ-system lemma to the finite reducts arising in global configurations. Altogether, this makes computing the invariants tractable.

2.2 Preliminaries

Notions from Classification Theory

Definition 2.2.1. We recall the following definitions, introduced in [She90].

1. A *cdt-pattern of height* κ is a sequence of formulas $\varphi_i(x; y_i)$ ($i < \kappa, i$ successor) and numbers $n_i < \omega$, and a tree of tuples $(a_\eta)_{\eta \in \omega^{<\kappa}}$ for which

    a) $p_\eta = \{\varphi_i(x; a_{\eta|i}) : i \text{ successor }, i < \kappa\}$ is consistent for $\eta \in \omega^\kappa$

    b) $\{\varphi_i(x; a_{\eta \frown \langle \alpha \rangle}) : \alpha < \omega, i = l(\eta) + 1\}$ is n_i-inconsistent.

2. An *inp-pattern of height* κ is a sequence of formulas $\varphi_i(x; y_i)$ ($i < \kappa$), sequences $(a_{i,\alpha} : \alpha < \omega)$, and numbers $n_i < \omega$ so that

    a) for any $\eta \in \omega^\kappa$, $\{\varphi_i(x; a_{i,\eta(i)}) : i < \kappa\}$ is consistent.

    b) For any $i < \kappa$, $\{\varphi_i(x; a_{i,\alpha}) : \alpha < \omega\}$ is n_i-inconsistent.

3. An *sct-pattern of height* κ is a sequence of formulas $\varphi_i(x; y_i)$ $(i < \kappa)$ and a tree of tuples $(a_\eta)_{\eta \in \omega^{<\kappa}}$ so that

   a) For every $\eta \in \omega^\kappa$, $\{\varphi_\alpha(x; a_{\eta|\alpha}) : 0 < \alpha < \kappa, \alpha \text{ successor}\}$ is consistent.

   b) If $\eta \in \omega^\alpha$, $\nu \in \omega^\beta$, α, β are successors, and $\nu \perp \eta$ then $\{\varphi_\alpha(x; a_\eta), \varphi_\beta(x; a_\nu)\}$ are inconsistent.

4. For $X \in \{\mathrm{cdt}, \mathrm{sct}, \mathrm{inp}\}$, we define $\kappa^n_X(T)$ be the first cardinal κ so that there is no X-pattern of height κ in n free variables. We define $\kappa_X(T) = \sup\{\kappa^n_X\}$.

Example 2.2.2. Fix a regular uncountable cardinal κ and let $L = \langle E_\alpha : \alpha < \kappa \rangle$ be a language consisting of κ many binary relations. Let T_{sct} be the model companion of the L-theory asserting that each E_α is an equivalence relation and $\alpha < \beta$ implies E_β refines E_α. Let T_{inp} be the model companion of the L-theory which only asserts that each E_α is an equivalence relation. Now $\kappa_{\mathrm{cdt}}(T_{\mathrm{sct}}) = \kappa_{\mathrm{cdt}}(T_{\mathrm{sct}}) = \kappa^+$, and further $\kappa_{\mathrm{sct}}(T_{\mathrm{sct}}) = \kappa_{\mathrm{inp}}(T_{\mathrm{inp}}) = \kappa^+$. However, we have $\kappa_{\mathrm{inp}}(T_{\mathrm{sct}}) = \aleph_0$ and $\kappa_{\mathrm{sct}}(T_{\mathrm{inp}}) = \aleph_1$.

In order to simplify many of the arguments below, it will be useful to work with indiscernible trees and arrays. Define a language $L_{s,\lambda} = \{\triangleleft, \wedge, <_{lex}, P_\alpha : \alpha < \lambda\}$ where λ is a cardinal. We may view the tree $\kappa^{<\lambda}$ as an $L_{s,\lambda}$-structure in a natural way, interpreting $\triangleleft$ as the tree partial order, $\wedge$ as the binary meet function, $<_{lex}$ as the lexicographic order, and P_α as a predicate which identifies the αth level. See Chapter 1 for a detailed treatment.

Definition 2.2.3.

1. We say $(a_\eta)_{\eta \in \kappa^{<\lambda}}$ is an *s-indiscernible tree over* A if

$$\mathrm{qftp}_{L_{s,\lambda}}(\eta_0, \ldots, \eta_{n-1}) = \mathrm{qftp}_{L_{s,\lambda}}(\nu_0, \ldots, \nu_{n-1})$$

   implies $\mathrm{tp}(a_{\eta_0}, \ldots, a_{\eta_{n-1}}/A) = \mathrm{tp}(a_{\nu_0}, \ldots, a_{\nu_{n-1}}/A)$.

2. We say $(a_{\alpha,i})_{\alpha<\kappa, i<\omega}$ is a *mutually indiscernible array* over A if, for all $\alpha < \kappa$, $(a_{\alpha,i})_{i<\omega}$ is a sequence indiscernible over $A \cup \{a_{\beta,j} : \beta < \kappa, \beta \neq \alpha, j < \omega\}$.

The following facts are also in Chapter 1:

Fact 2.2.4. Let $(a_\eta : \eta \in \kappa^{<\lambda})$ be a tree s-indiscernible over a set of parameters C.

1. All paths have the same type over C: for any $\alpha, \nu \in \kappa^\lambda$, $\mathrm{tp}((a_{\eta|\alpha})_{\alpha<\lambda}/C) = \mathrm{tp}((a_{\nu|\alpha})_{\alpha<\lambda}/C)$.

2. Suppose $\{\eta_\alpha : \alpha < \gamma\} \subseteq \kappa^{<\lambda}$ satisfies $\eta_\alpha \perp \eta_{\alpha'}$ whenever $\alpha \neq \alpha'$. Then the array $(b_{\alpha,\beta})_{\alpha<\gamma, \beta<\kappa}$ defined by

$$b_{\alpha,\beta} = a_{\eta_\alpha \frown \langle \beta \rangle}$$

   is mutually indiscernible over C.

Fact 2.2.5. 1. If there is an inp-pattern of height κ modulo T, then there is an inp-pattern $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$, $(a_{\alpha,i})_{\alpha<\kappa,i<\omega}$ so that $(a_{\alpha,i})_{\alpha<\kappa,i<\omega}$ is a mutually indiscernible array.

2. If there is an sct-pattern (cdt-pattern) of height κ modulo T, then there is an sct-pattern (cdt-pattern) $\varphi_\alpha(x; y_\alpha)$, $(a_\eta)_{\eta\in\omega^{<\kappa}}$ so that $(a_\eta)_{\eta\in\omega^{<\kappa}}$ is an s-indiscernible tree.

Fraïssé Theory

We will recall some basic facts from Fraïssé theory, from [Hod93]. Let L be a finite language and let $\mathbb{K}$ be a non-empty finite or countable set of finitely generated L-structures which has HP, JEP, and AP. Such a class $\mathbb{K}$ is called a *Fraïssé class.* Then there is an L-structure D, unique up to isomorphism, such that D has cardinality $\leq \aleph_0$, $\mathbb{K}$ is the age of D, and D is ultrahomogeneous. We call D the *Fraïssé limit* of $\mathbb{K}$, which we sometimes denote $\mathrm{Flim}(\mathbb{K})$. We say that $\mathbb{K}$ is *uniformly locally finite* if there is a function $g : \omega \to \omega$ so that a structure in $\mathbb{K}$ generated by n elements has cardinality at most $g(n)$. If $\mathbb{K}$ is a countable uniformly locally finite set of finitely generated L-structures and $T = \mathrm{Th}(D)$, then T is $\aleph_0$-categorical and has quantifier-elimination.

Fact 2.2.6. [KPT05] Let A be a countable structure. Then A is ultrahomogeneous if and only if it satisfies the following extension property: if B, C are finitely generated and can be embedded into A, $f : B \to A$, $g : B \to C$ are embeddings then there is an embedding $h : C \to A$ so that $h \circ g = f$.

The following is a straight-forward generalization of [KPT05, Proposition 5.2]:

Lemma 2.2.7. *Suppose $L \subseteq L'$, and $\mathbb{K}$ is a Fraïssé class of L-structures and $\mathbb{K}'$ is a Fraïssé class of L'-structures satisfying the following two conditions:*

1. *$A \in \mathbb{K}$ if and only if there is a $D' \in \mathbb{K}'$ so that A is an L-substructure of $D' \restriction L$.*

2. *If $A, B \in \mathbb{K}$, $\pi : A \to B$ is an L-embedding, and $C \in \mathbb{K}'$ with $C = \langle A\rangle^C_{L'}$, then there is a $D \in \mathbb{K}'$ with $D = \langle B\rangle^D_{L'}$ and an L'-embedding $\tilde{\pi} : C \to D$ extending π.*

Then $\mathrm{Flim}(\mathbb{K}') \restriction L = \mathrm{Flim}(\mathbb{K})$.

Proof. Let $F' = \mathrm{Flim}(\mathbb{K}')$ and suppose $F = F' \restriction L$. Fix $A_0, B_0 \in \mathbb{K}$ and an L-embedding $\pi : A_0 \to B_0$. Suppose $\varphi : A_0 \to F$ is an L-embedding. Let $E = \langle\varphi(A_0)\rangle^{F'}_{L'}$. Up to isomorphism over A_0, there is a unique $C \in \mathbb{K}'$ containing A_0 so that $C = \langle A_0\rangle^C_{L'}$ and $\tilde{\varphi} : C \to F'$ is an L'-embedding extending φ and with $E = \tilde{\varphi}(C)$. By hypothesis, there is some $D \in \mathbb{K}'$ with $B_0 \subseteq D$ and $D = \langle B_0\rangle^D_{L'}$ and there is an L'-embedding $\tilde{\pi} : C \to D$ extending π. By the extension property for F', there is an L'-embedding $\psi : D \to F'$ so that $\psi \circ \tilde{\pi} = \tilde{\varphi}$ and hence $\psi \circ \pi = \varphi$. As $\psi \restriction B_0$ is an L-embedding, this shows the extension property for F. So F is ultrahomogeneous, and $\mathrm{Age}(F) = \mathbb{K}$ by (1) so $F \cong \mathrm{Flim}(\mathbb{K})$, which completes the proof. □

Strong Colorings

Definition 2.2.8. [She94] Write $\mathrm{Pr}_1(\lambda, \mu, \theta, \chi)$ for the assertion: there is a coloring $c : [\lambda]^2 \to \theta$ so that for any $A \subseteq [\lambda]^{<\chi}$ of size μ consisting of pairwise disjoint subsets of λ and any color $\gamma < \theta$ there are $a, b \in A$ with $\max(a) < \min(b)$ with $c(\{\alpha, \beta\}) = \gamma$ for all $\alpha \in a$, $\beta \in b$.

Note, for example, that $\mathrm{Pr}_1(\lambda, \lambda, 2, 2)$ holds if and only if $\lambda \not\to (\lambda)^2_2$ - i.e. λ is not weakly compact.

Observation 2.2.9. For fixed λ, if $\mu \leq \mu'$, $\theta' \leq \theta$, $\chi' \leq \chi$, then

$$\mathrm{Pr}_1(\lambda, \mu, \theta, \chi) \implies \mathrm{Pr}_1(\lambda, \mu', \theta', \chi').$$

Proof. Fix $c : [\lambda]^2 \to \theta$ witnessing $\mathrm{Pr}_1(\lambda, \mu, \theta, \chi)$. Define a new coloring $c' : [\lambda]^2 \to \theta'$ by $c'(\{\alpha, \beta\}) = c(\{\alpha, \beta\})$ if $c(\{\alpha, \beta\}) < \theta'$ and $c'(\{\alpha, \beta\}) = 0$ otherwise. Now suppose $A \subseteq [\lambda]^{<\chi'}$ is a family of pairwise disjoint sets with $|A| \geq \mu'$. Then, in particular, $A \subseteq [\lambda]^{<\chi}$ and $|A| \geq \mu$ so for any $\gamma < \theta'$, as $\gamma < \theta$, there are $a, b \in A$ with $\max(a) < \min(b)$ with $c'(\{\alpha, \beta\}) = c(\{\alpha, \beta\}) = \gamma$ for all $\alpha \in a$, $\beta \in b$, using $\mathrm{Pr}_1(\lambda, \mu, \theta, \chi)$ and the definition of c'. This shows c' witnesses $\mathrm{Pr}_1(\lambda, \mu', \theta', \chi')$. □

Galvin proved Pr_1 holds in some form for arbitrary successor cardinals from instances of GCH. Considerably later, Shelah proved that Pr_1 holds in a strong form for the double-successors of arbitrary regular cardinals in ZFC.

Fact 2.2.10. [Gal80] If λ is an infinite cardinal and $2^\lambda = \lambda^+$, then $\mathrm{Pr}_1(\lambda^+, \lambda^+, \lambda^+, \aleph_0)$.

Fact 2.2.11. [She97] The principle $\mathrm{Pr}_1(\lambda^{++}, \lambda^{++}, \lambda^{++}, \lambda)$ holds for every regular cardinal λ.

2.3 The main construction

From strong colorings, we construct theories with $\kappa_{\mathrm{sct}}(T) + \kappa_{\mathrm{inp}}(T) < \kappa_{\mathrm{cdt}}(T)$. For each regular uncountable cardinal κ and coloring $f : [\kappa]^2 \to 2$ we build a theory $T^*_{\kappa,f}$ which comes equipped with a canonical cdt-pattern of height κ, in which the consistency of two incomparable nodes, one on level α and another on level β, is determined by the value of the coloring $f(\{\alpha, \beta\})$. We then analyze the possible inp- and sct-patterns that arise in models of $T^*_{\kappa,f}$ and show that the existence of an inp- or sct-pattern of height κ implies certain homogeneity properties about the coloring f. If there is an inp-pattern of height κ, we can conclude that f has a homogeneous set of size κ. In the case that there is an sct-pattern of height κ, we cannot quite get a homogeneous set, but one nearly so: we prove in this case that there is a color $\gamma \in \{0, 1\}$ and a collection $(v_\alpha : \alpha < \kappa)$ of pairwise disjoint finite subsets of κ so that given $\alpha < \alpha'$, there are $\xi \in v_\alpha, \zeta \in v_{\alpha'}$ so that $f(\{\xi, \zeta\}) = \gamma$. This is precisely the kind of homogeneity which a strong coloring witnessing $\mathrm{Pr}_1(\kappa, \kappa, 2, \aleph_0)$ explicitly prohibits. The theory associated to such a coloring, then, gives the desired counterexample.

Building a Theory

Suppose κ is a regular uncountable cardinal. We define a language $L_\kappa = \langle O, P_\alpha, f_{\alpha\beta}, p_\alpha : \alpha \leq \beta < \kappa \rangle$, where O and all the P_α are unary predicates and the $f_{\alpha\beta}$ and p_α are unary functions. Given a subset $w \subseteq \kappa$, let $L_w = \langle O, P_\alpha, f_{\alpha\beta}, p_\alpha : \alpha \leq \beta, \alpha, \beta \in w \rangle$. Given a function $f : [\kappa]^2 \to 2$, we define a universal theory $T_{\kappa,f}$ with the following axiom schemas:

1. O and the P_α are pairwise disjoint;
2. $f_{\alpha\alpha}$ is the identity function, for all $\alpha < \beta$,

$$(\forall x) \left[(x \notin P_\beta \to f_{\alpha\beta}(x) = x) \wedge (x \in P_\beta \to f_{\alpha\beta}(x) \in P_\alpha) \right],$$

and if $\alpha < \beta < \gamma$, then

$$(\forall x \in P_\gamma)[f_{\alpha\gamma}(x) = (f_{\alpha\beta} \circ f_{\beta\gamma})(x)].$$

3. For all $\alpha < \kappa$,

$$(\forall x) \left[(x \notin O \to p_\alpha(x) = x) \wedge (p_\alpha(x) \neq x \to p_\alpha(x) \in P_\alpha) \right].$$

4. If $\alpha < \beta < \kappa$ and $f(\{\alpha, \beta\}) = 0$, then we have the axiom $(\forall z \in O)[p_\alpha(z) \neq z \wedge p_\beta(z) \neq z \to p_\alpha(z) = (f_{\alpha\beta} \circ p_\beta)(z)]$.

The O is for “objects" and $\bigcup P_\alpha$ is a tree of “parameters” where each P_α names nodes of level α. The functions $f_{\alpha\beta}$ map elements of the tree at level β to their unique predecessor at level α. So the tree partial order is coded in a highly non-uniform way, for each pair of levels. The p_α's should be considered as partial functions on O which connect objects to elements of the tree. Axiom (4) says, in essence, that if $f(\{\alpha, \beta\}) = 0$, then the only way for an object to connect to a node on level α and a node on level β is if these two nodes lie along a path in the tree.

Lemma 2.3.1. *Define a class of finite structures*

$$\mathbb{K}_w = \{ \text{ finite models of } T_{\kappa,f} \upharpoonright L_w \}.$$

Then for finite w, $\mathbb{K}_w$ is a Fraïssé class and, moreover, it is uniformly locally finite.

Proof. The axioms for $T_{\kappa,f}$ are universal so HP is clear. JEP and AP are proved similarly, so we will give the argument for AP only. Suppose A includes into B and C where $A, B, C \in \mathbb{K}_w$ and $B \cap C = A$. $B \cup C$ may be viewed as an L_w-structure by interpreting each predicate Q of L_w so that $Q^{B \cup C} = Q^B \cup Q^C$ and similarly interpreting $g^{B \cup C} = g^B \cup g^C$ for all the function symbols $g \in L_w$. It is easy to check that $B \cup C$ is a model of $T_{\kappa,f} \upharpoonright L_w$. To see uniform local finiteness, just observe that a set of size n can generate a model of size at most $(|w| + 1)n$ in virtue of the way that the functions are defined. $\square$

Hence, for each finite $w \subset \kappa$, there is a countable ultrahomogeneous L_w-structure M_w with $\mathrm{Age}(M_w) = \mathbb{K}_w$. Let $T^*_w = \mathrm{Th}(M_w)$.

Lemma 2.3.2. *Suppose $w \subseteq v \subseteq \kappa$ and v, w are both finite. Then $T^*_w \subset T^*_v$.*

Proof. Suppose w is a finite subset of κ and $\gamma \in \kappa \setminus w$. By induction, it suffices to consider the case when $v = w \cup \{\gamma\}$. We will use Lemma 2.2.7. It is clear that if $D \in \mathbb{K}_w$ and A is an L_v-substructure of $D|_{L_v}$ then $A \in K_v$. For the other direction, suppose $A \in \mathbb{K}_v$ and we will construct $D \in \mathbb{K}_w$ so that A is an L_v-substructure of $D|_{L_v}$. If γ is the greatest element of v, let D be the L_w-structure expanding A with $P^D_\gamma = \emptyset$ and p_γ and each $f_{\beta\gamma}$ interpreted as the identity functions for all $\beta \in w$. It is easy to check that $D \in \mathbb{K}_w$. If γ is not the greatest element of w, let α^* be the least element of w greater than γ. Let D have underlying set $A \cup \{*_d : d \in P^A_{\alpha^*}\}$, where the $*_d$ denote new formal elements. Interpret the predicates of L_w on D so that $P^D_\gamma = \{*_d : d \in P^A_{\alpha^*}\}$, $O^D = O^A$ and $P^D_\beta = P^A_\beta$ for all $\beta \in v$. If $\beta \leq \delta$ are in v, interpret $f^D_{\beta\delta}|_{P^D_\delta} = f^A_{\beta\delta}|_{P^A_\delta}$ and to be the identity function elsewhere. If $\beta \in v$ and $\beta > \gamma$ and $c \in P^D_\beta$, define $f^D_{\gamma\beta}(c) = *_{f_{\alpha^*\beta}(c)}$ and to be the identity elsewhere. If $\beta \in v$ and $\beta < \gamma$ then define $f^D_{\beta\gamma}$ to be the identity outside of P^D_γ and so, on P^D_γ, $f_{\beta\gamma}(*_d) = f_{\beta\alpha_*}(d)$. Finally, interpret $f^D_{\gamma\gamma}$ to be the identity function. It is clear from the construction that A is an L_v-substructure of $D|_{L_v}$ and easy to check that $D \in \mathbb{K}_w$. This shows that the condition of Lemma 2.2.7(1) is satisfied.

To verify the condition of Lemma 2.2.7(2), must show that if $A, B \in \mathbb{K}_w$, $\pi : A \to B$ is an L_w-embedding, and $C \in \mathbb{K}_v$ with $C = \langle A \rangle^C_{L_v}$, then there is some $D \in \mathbb{K}_v$ with $\langle B \rangle^D_{L_v} = D$ and an L_v-embedding $\tilde{\pi} : C \to D$ extending π.

Let A, B, π and C as above be given, and we will construct D. We may assume that $C \cap B = \emptyset$. The requirement that $C = \langle A \rangle^C_{L_v}$ implies that the only elements of the underlying set of C that are not already in A appear in $P_\gamma(C)$. As a set, define D to be the set

$$D = B \cup P_\gamma(C) \cup \{*_d : d \in P_{\alpha^*}(B) \setminus \pi(P_{\alpha^*}(C))\}$$

where α^* is the least element of w greater than γ (and if there is no such, then let $D = B \cup P_\gamma(C)$) and, as above, the $*_d$ denote new formal elements. Let $\tilde{\pi}$ be the map extending π which is the identity on $P_\gamma(C)$.

Let $P^D_\gamma = P_\gamma(C) \cup \{*_d : d \in P_{\alpha^*}(B) \setminus \pi(P_{\alpha^*}(C))\}$ and interpret all other unary predicates on D to coincide with their interpretation on B. Define the interpretations of the functions on D as follows: for any pair $\alpha, \beta \in w$ with $\alpha \leq \beta$, and $b \in B$ let $f^D_{\alpha\beta}(b) = f^B_{\alpha\beta}(b)$. Interpret $f^D_{\gamma\alpha^*}$ on $P_\gamma(D)$ so that $c \in P_\gamma(C)$ then $f^D_{\gamma\alpha^*}(\pi(c)) = f^C_{\gamma\alpha^*}(c) = \tilde{\pi}(f^C_{\gamma\alpha^*}(c))$ and so that $f_{\gamma\alpha^*}(d) = *_d$. Then let $f^D_{\gamma\alpha^*}$ be the identity outside of $P_{\alpha^*}(D)$. Now the interpretations of the others are forced: if $\alpha \geq \alpha^*$ and $\beta \leq \beta^*$, define $f^D_{\gamma\alpha}$ so that $f^D_{\gamma\alpha} = f^B_{\gamma\alpha^*} \circ f^D_{\alpha^*\alpha}$ and $f^D_{\beta\gamma} = f^B_{\beta\beta^*} \circ f^D_{\beta^*\gamma}$. Finally, interpret $p^D_\alpha = p^B_\alpha$ for $\alpha \in w$ and let p^D_γ be interpreted so that

$$p^D_\gamma(d) = \begin{cases} p^C_\gamma(c) = \tilde{\pi}(p^C_\gamma(c)) & \text{if } d = \pi(c), c \in O^C \\ d & \text{otherwise.} \end{cases}$$

By construction $p^D_\gamma(x) = x$ for all x outside of $O(D)$ and $p^D_\gamma(x) \neq x$ implies $p^D_\gamma(x) = p^C(y) \in P_\gamma(C) = P_\gamma(D)$ for some $y \in C$ with $\pi(c) = x$, so D satisfies axiom (3). It is clear that $\tilde{\pi} : C \to D$ is a function which extends π and, moreover, it preserves all L_v-structure so $\tilde{\pi}$ is an L_v-embedding. Finally, we check axiom (4). Suppose there are $\alpha < \beta$ with $\alpha, \beta \in v$ and $f(\{\alpha, \beta\}) = 0$, together with $d \in O(D)$, $g \in P_\alpha(D)$, $h \in P_\beta(D)$ so that $p_\alpha(d) = g$ and $p_\beta(d) = h$. We must show that $f_{\alpha\beta}(h) = g$. If $\alpha, \beta \in w$, then this is immediate, as axiom (4) is satisfied in B. So consider the case that $\beta = \gamma$. Now $p_\gamma(d) = h \neq d$ implies, by construction, that there is some $c \in C$ so that $d = \pi(c)$. As $p_\alpha(d) \in P_\alpha(D)$ and $p_\beta(d) \in P_\beta(D)$, it follows that $p_\alpha(c) \in P_\alpha(C)$ and $p_\beta(c) \in P_\beta(C)$. As axiom (4) is satisfied in C, we know that $(f_{\alpha\beta} \circ p_\beta)(c) = p_\alpha(c)$. It follows that $(f_{\alpha\beta} \circ p_\beta)(d) = p_\alpha(d)$ in D, which shows $f_{\alpha\beta}(h) = g$ in D. The case that $\alpha = \gamma$ is entirely similar, so axiom (4) is satisfied in D. □

Lemma 2.3.2 shows that we may construct a complete L_κ-theory as a union of the theories of the Fraïssé limits of the $\mathbb{K}_w$ for $w \subset \kappa$ finite. Define the theory $T^*_{\kappa,f}$ as the union of the T^*_w for all finite $w \subset \kappa$. Because each T^*_w is complete and eliminates quantifiers, it follows that $T^*_{\kappa,f}$ is a complete theory extending $T_{\kappa,f}$ which eliminates quantifiers.

2.4 Analysis of the invariants

For this section, we will fix κ a regular uncountable cardinal, a coloring $f : [\kappa]^2 \to 2$, and a monster model $\mathbb{M} \models T^*_{\kappa,f}$.

Rectification and other preparatory lemmas

Definition 2.4.1. Given $X \in \{\text{inp}, \text{sct}\}$, we say an X-pattern $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$ of height κ is *rectified* if the following conditions hold:

1. There is a sequence of sets $(w_\alpha : \alpha < \kappa)$ach $\varphi_\alpha(x; y_\alpha)$ is contained in the language L_{w_α} where $w_\alpha \subset \kappa$ is finite. The sets $(w_\alpha : \alpha < \kappa)$ form a Δ-system with root r so that for all $\alpha < \kappa$, $\max r < \min(w_\alpha \setminus r)$ and if $\alpha < \alpha'$, $\max(w_\alpha \setminus r) < \min(w_{\alpha'} \setminus r)$. Moreover, each w_α has the same cardinality.

2. The witnessing parameters are indiscernible in the relevant sense (s-indiscernible in the case that $X = \text{sct}$, mutually indiscernible in the case of $X = \text{inp}$).

3. $\varphi_\alpha(x; y_\alpha)$ isolates the L_{w_α}-type of x over y_α.

4. The tuples in the witnessing parameters are closed under the functions in the language corresponding to their level: if $X = \text{inp}$ and $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$ is an inp-pattern witnessed by the mutually indiscernible array $(a_{\alpha,i})_{\alpha<\kappa,i<\omega}$ then for all $\alpha < \kappa$ and $i < \omega$, $a_{\alpha,i}$ is closed under the functions of L_{w_α}. Similarly, if $X = \text{sct}$ and $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$

is an sct-pattern witnessed by $(a_\eta)_{\eta\in\omega^{<\kappa}}$, then for all $\eta \in \omega^{<\kappa}$, the tuple a_η is closed under the functions of $L_{w_{l(\eta)}}$.

By the *associated* Δ*-system* of a rectified X-pattern, we mean the Δ-system $(w_\alpha : \alpha < \kappa)$ so that $\varphi_\alpha \in L_{w_\alpha}$. We will consistently denote the root $r = \{\alpha_i : i < n\}$ and the sets $v_\alpha = w_\alpha \setminus r = \{\beta_{\alpha,i} : i < m\}$ where the enumerations are increasing.

Lemma 2.4.2. *Given $X \in \{inp, sct\}$, if there is an X-pattern of height κ in T, there is a rectified one.*

Proof. Given an X-pattern $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$ one can choose some finite $w_\alpha \subset \kappa$ so that $\varphi_\alpha(x; y_\alpha)$ is in L_{w_α}. Apply the Δ-system lemma to the collection $(w_\alpha : \alpha < \kappa)$ to find some $I \subseteq \kappa$ so that $(w_\alpha : \alpha \in I)$ forms a Δ-system with root r. By pigeonhole and the regularity of κ, we may assume $|w_\alpha| = m$ for all α, for all $\alpha < \kappa$ $\max r < \min(w_\alpha \setminus r)$, and if $\alpha < \alpha'$, $\max(w_\alpha \setminus r) < \min(w_{\alpha'} \setminus r)$. By renaming, we may assume $I = \kappa$.

If $X =$ inp, we may take the parameters witnessing that $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$ to be a mutually indiscernible array $(a_{\alpha,i})_{\alpha<\kappa,i<\omega}$ by Fact 2.2.5(1). Moreover, mutual indiscernibility is clearly preserved after replacing each $a_{\alpha,i}$ by its closure under the functions of L_{w_α} and by $\aleph_0$-categoricity of $T^*_{w_\alpha}$ this stays finite. Let $b \models \{\varphi_\alpha(x; a_{\alpha,0}) : \alpha < \kappa\}$. Using again the $\aleph_0$-categoricity of $T^*_{L_{w_\alpha}}$, replace $\varphi_\alpha(x; y_\alpha)$ by a complete L_{w_α}-formula $\varphi'_\alpha(x; y_\alpha)$ so that $\varphi'_\alpha(x; a_{\alpha,0})$ isolates the type $\mathrm{tp}_{L_{w_\alpha}}(b/a_{\alpha,0})$. By mutual indiscernibility, if $f : \kappa \to \omega$ is a function, there is $\sigma \in \mathrm{Aut}(\mathbb{M})$ so that $\sigma(a_{\alpha,0}) = a_{\alpha,f(\alpha)}$ for all $\alpha < \kappa$. Then $\sigma(b) \models \{\varphi'_\alpha(x; a_{\alpha,f(\alpha)}) : \alpha < \kappa\}$ so paths are consistent. The row-wise inconsistency is clear so $(\varphi'_\alpha(x; y_\alpha) : \alpha < \kappa)$ forms an inp-pattern.

If $X =$ sct, we may take the parameters witnessing that $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$ is an sct-pattern to be s-indiscernible, by Fact 2.2.5(2). s-indiscernibility is preserved by replacing each a_η by its closure under the functions of $L_{w_{l(\eta)}}$ and this closure is finite. Let $b \models \{\varphi_\alpha(x; a_{0^\alpha}) : \alpha < \kappa\}$ and replace φ_α by $\varphi'_\alpha(x; y_\alpha)$, a complete L_{w_α}-formula isolating $tp_{L_{w_\alpha}}(b/a_{0^\alpha})$. For all $\eta \in \omega^\kappa$, there is a $\sigma \in \mathrm{Aut}(\mathbb{M})$ so that $\sigma(a_{0^\alpha}) = a_{\eta|\alpha}$. Then $\sigma(b) \models \{\varphi'_\alpha(x; a_{\eta|\alpha}) : \alpha < \kappa\}$ so paths are consistent. Incomparable nodes remain inconsistent, so $(\varphi'_\alpha(x; y_\alpha) : \alpha < \kappa)$ forms an sct-pattern. □

Computing κ_{cdt}

Proposition 2.4.3. $\kappa_{cdt}(T^*_{\kappa,f}) = \kappa^+$.

Proof. First, we will show $\kappa_{\mathrm{cdt}}(T^*_{\kappa,f}) \geq \kappa^+$. We will construct a cdt-pattern of height κ. By recursion on $\alpha < \kappa$, we will construct a tree of tuples $(a_\eta)_{\eta\in\omega^{<\kappa}}$ so that $l(\eta) = \beta$ implies $a_\eta \in P_\beta$ and if $\eta \trianglelefteq \nu$ with $l(\eta) = \beta$ and $l(\nu) = \gamma$, then $f_{\beta\gamma}(a_\nu) = a_\eta$. For $\alpha = 0$, choose an arbitrary $a \in P_0$ and let $a_\emptyset = a$. Now suppose given $(a_\eta)_{\eta\in\omega^{\leq\alpha}}$. For each $\eta \in \omega^\alpha$, choose an infinite set $\{b_i : i < \omega\} \subseteq f^{-1}_{\alpha\alpha+1}(a_\eta)$. Define $a_{\eta\frown\langle i\rangle} = b_i$. This gives us $(a_\eta)_{\eta\in\omega^{\leq\alpha+1}}$ with the desired properties. Now suppose δ is a limit and we've defined $(a_\eta)_{\eta\in\omega^{\leq\alpha}}$ for all $\alpha < \delta$. Given any $\eta \in \omega^\delta$, we may, by saturation, find an element $b \in \bigcap_{\alpha<\delta} f^{-1}_{\alpha\delta}(a_{\eta|\alpha})$. Then we can set $a_\eta = b$. This gives $(a_\eta)_{\eta\in\omega^{\leq\delta}}$ and completes the construction.

Given $\alpha < \kappa$, let $\varphi_\alpha(x; y)$ be the formula $p_\alpha(x) = y$. For any $\eta \in \omega^\kappa$, $\{\varphi_\alpha(x; a_{\eta|\alpha}) : \alpha < \kappa\}$ is consistent and, for all $\nu \in \omega^{<\kappa}$, $\{\varphi_{l(\nu)+1}(x; a_{\nu\frown\langle i\rangle}) : i < \omega\}$ is 2-inconsistent. We have thus exhibited a cdt-pattern of height κ so $\kappa_{\mathrm{cdt}}(T^*_{\kappa,f}) \geq \kappa^+$.

Using quantifier-elimination, it is easy to check that each theory T^*_w is stable (in fact, ω-stable) for any finite $w \subset \kappa$. Hence $T^*_{\kappa,f}$ is stable and therefore $\kappa_{\mathrm{cdt}}(T^*_{\kappa,f}) \leq |T^*_{\kappa,f}|^+ = \kappa^+$. This yields the desired equality. $\square$

Case 1: $\kappa_{\mathbf{inp}} = \kappa^+$

Lemma 2.4.4. *Fix a collection of ordinals $< \kappa$ $(\beta_{\alpha,i})_{\alpha<\kappa,i<2}$ so that if $\alpha < \alpha' < \kappa$, then $\beta_{\alpha,0} \leq \beta_{\alpha,1}$, $\beta_{\alpha',0} \leq \beta_{\alpha',1}$, $\beta_{\alpha,0} \leq \beta_{\alpha',0}$ and $\beta_{\alpha,1} < \beta_{\alpha',1}$. Suppose that there is a mutually indiscernible array $(c_{\alpha,k})_{\alpha<\kappa,k<\omega}$ so that, with $\varphi_\alpha(x; y_\alpha)$ defined by $(f_{\beta_{\alpha,0}\beta_{\alpha,1}} \circ p_{\beta_{\alpha,1}})(x) = y_\alpha$, $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$, $(c_{\alpha,k})_{\alpha<\kappa,k<\omega}$ forms an inp-pattern of height κ. Then for all $\alpha < \alpha'$, $f(\{\beta_{\alpha,1}, \beta_{\alpha',1}\}) = 1$.*

Proof. If $\alpha < \alpha'$ and $f(\{\beta_{\alpha,1}, \beta_{\alpha',1}\}) = 0$, then $p_{\beta_{\alpha,1}}(x) = (f_{\beta_{\alpha,1}\beta_{\alpha',1}} \circ p_{\beta_{\alpha',1}})(x)$ for any x with $p_{\beta_{\alpha,1}}(x) \neq x$ and $p_{\beta_{\alpha',1}}(x) \neq x$, and hence

$$\begin{aligned}
(f_{\beta_{\alpha,0}\beta_{\alpha,1}} \circ p_{\beta_{\alpha,1}})(x) &= (f_{\beta_{\alpha,0}\beta_{\alpha,1}} \circ f_{\beta_{\alpha,1}\beta_{\alpha',1}} \circ p_{\beta_{\alpha',1}})(x) \\
&= (f_{\beta_{\alpha,0}\beta_{\alpha',1}} \circ p_{\beta_{\alpha',1}})(x) \\
&= (f_{\beta_{\alpha,0},\beta_{\alpha',0}} \circ f_{\beta_{\alpha',0}\beta_{\alpha',1}} \circ p_{\beta_{\alpha',1}})(x),
\end{aligned}$$

for any x with $p_{\beta_{\alpha,1}}(x) \neq x$ and $p_{\beta_{\alpha',1}}(x) \neq x$. Consequently,

$$\{(f_{\beta_{\alpha,0}\beta_{\alpha,1}} \circ p_{\beta_{\alpha,1}})(x) = c_{\alpha,k}, (f_{\beta_{\alpha',0}\beta_{\alpha',1}} \circ p_{\beta_{\alpha',1}})(x) = c_{\alpha',k'}\}$$

is consistent only if $c_{\alpha,k} = f_{\beta_{\alpha,0}\beta_{\alpha',0}}(c_{\alpha',k'})$, which contradicts the definition of inp-pattern. $\square$

For the remainder of this subsection, we fix a rectified inp-pattern $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$, $(a_{\alpha,i})_{\alpha<\kappa,i<\omega}$ and, by [Che14, Corollary 2.9], we may assume $l(x) = 1$. The associated Δ-system is denoted $(w_\alpha : \alpha < \kappa)$ with root $r = \{\alpha_i : i < n\}$ and $w_\alpha \setminus r = v_\alpha = \{\beta_{\alpha,j} : j < m\}$.

Lemma 2.4.5. *For all $\alpha < \kappa$, $\varphi_\alpha(x; y_\alpha) \vdash x \in O$.*

Proof. First, note that we may assume that there is a predicate $Q \in \{O, P_{\alpha_i} : i < n\}$ so that $\varphi_\alpha(x; y_\alpha) \vdash x \in Q$ for all $\alpha < \kappa$ – if not, using that the w_α's form a Δ-system, there would be some $\alpha < \kappa$ so that $\varphi_\alpha(x; y_\alpha)$ implies that x is not contained in any predicate of L_{w_α} in which case it is easy to check that $\{\varphi_\alpha(x; a_{\alpha,i}) : i < \omega\}$ is consistent, contradicting the definition of inp-pattern. So we must show that $\varphi_\alpha(x; y_\alpha) \vdash P_{\alpha_i}$ for some $i < n$ is impossible. Write each tuple in the array $a_{\alpha,i}$ as $a_{\alpha,i} = (b_{\alpha,i}, c_{\alpha,i}, d_{\alpha,i}, e_{\alpha,i})$ where the elements of $b_{\alpha,i}$ are in O, the elements of $c_{\alpha,i}$ are in predicates indexed by the root $\bigcup_{i<n} P_{\alpha_i}$, the elements of $d_{\alpha,i}$ are in predicates outside the root $\bigcup_{j<m} P_{\beta_{\alpha,j}}$, and the elements of $e_{\alpha,i}$ are in any predicate of L_{w_α}. By quantifier-elimination, each $\varphi_\alpha(x; a_{\alpha,i})$ may be expressed as a conjunction of

1. $x \in P_{\alpha_i}$

2. $x \neq (a_{\alpha,i})_l$ for all $l < l(a_{\alpha,i})$

3. $(f_{\gamma\alpha_i}(x) = (c)_l)^{t_{\gamma,l}}$ for all $l < l(c_{\alpha,i})$ and $\gamma \in w_\alpha$ less than α_i and some $t_{\gamma,l} \in \{0,1\}$.

For each $k < i$, let γ_k be the least ordinal $< \kappa$ so that $\varphi_{\gamma_k}(x; a_{\gamma_k,0}) \vdash f_{\alpha_k\alpha_i}(x) = c$ for some $c \in c_{\gamma_k,0}$ and 0 if there is no such. Let $\gamma = \max\{\gamma_k : k < i\}$. We claim that $\{\varphi_{\gamma+1}(x; a_{\gamma+1,j}) : j < \omega\}$ is consistent. For all $j < \omega$,

$$\{\varphi_\alpha(x; a_{\alpha,0}) : \alpha \leq \gamma\} \cup \{\varphi_{\gamma+1}(x; a_{\gamma+1,j})\}$$

is consistent so any equality of the form $f_{\alpha_k\alpha_i}(x) = c$ implied by $\varphi_{\gamma+1}(x; a_{\gamma+1,j})$ is already implied by $\varphi_\alpha(x; a_{\alpha,0})$ by our choice of γ and any inequality of the form $f_{\alpha_k\alpha_i}(x) \neq c$ implied by $\varphi_{\gamma+1}(x; a_{\gamma+1,j})$ is compatible with $\{\varphi_\alpha(x; a_{\alpha,0}) : \alpha \leq \gamma\}$. Choosing a realization $b \models \{\varphi_\alpha(x; a_{\alpha,0}) : \alpha \leq \gamma\}$ satisfying every inequality of the form $f_{\alpha_k\alpha_i}(x) \neq c$ implied by the $\varphi_{\gamma+1}(x; a_{\gamma+1,j})$ yields a realization of $\{\varphi_{\gamma+1}(x; a_{\gamma+1,j}) : j < \omega\}$. This contradicts the definition of inp-pattern. □

Proposition 2.4.6. *There is a subset $H \subseteq \kappa$ with $|H| = \kappa$ so that f is constant on $[H]^2$.*

Proof. By quantifier elimination and Lemma 2.4.5, for each $\alpha < \kappa$, $\varphi_\alpha(x; a_{\alpha,0})$ is a conjunction of the following:

1. $x \in O$

2. $x \neq (a)_l$ for all $l < l(a)$

3. $(p_\gamma(x) = x)^t$.

4. The values of the p_γ and how they descend in the tree:

   a) $((f_{\delta\gamma} \circ p_\gamma)(x) = (a_{0,\alpha})_l)^t$ for $l < l(a_{0,\alpha})$, $\delta \leq \gamma$ in w_α.

   b) $((f_{\delta\gamma} \circ p_\gamma)(x) = (f_{\delta\gamma'} \circ p_{\gamma'})(x))^t$ for $\delta, \gamma, \gamma' \in w_\alpha$ with $\delta \leq \gamma < \gamma'$.

Let $\gamma < \kappa$ be some ordinal so that for any $\alpha < \kappa$ if there is a $c \in a_{\alpha,0}$ so that $\varphi_\alpha(x; a_{\alpha,0}) \vdash (f_{\alpha_i\alpha_{i'}} \circ p_{\alpha_{i'}})(x) = c$ for some $i \leq i' < n$, then there is some $\alpha' < \gamma$ so that $\varphi_{\alpha'}(x; a_{\alpha',0}) \vdash (f_{\alpha_i\alpha_{i'}} \circ p_{\alpha_{i'}})(x) = c$. As the root $r = \{\alpha_i : i < n\}$ is finite and the all 0's path is consistent, such a γ must exist.

Claim: Given $\alpha < \kappa$, there are $\epsilon_\alpha \leq \epsilon'_\alpha \in w_\alpha$ and pairwise distinct $c_{\alpha,k} \in a_{\alpha,k}$ so that $\varphi_\alpha(x; a_{\alpha,k}) \vdash (f_{\epsilon_\alpha\epsilon'_\alpha} \circ p_{\epsilon'})(x) = c_{\alpha,k}$.

Proof of claim: Suppose not. Then, by the description of $\varphi_\alpha(x; a_{\alpha,k})$ given above, the partial type

$$\{\varphi_\alpha(x; a_{\alpha,k}) : k < \omega\}$$

is equivalent to a finite number of equations common to each instance $\varphi_\alpha(x; a_{0,k})$ and an infinite collection of inequations. Then, it is easy to see then that $\{\varphi_\alpha(x; a_{0,k}) : k < \omega\}$ is consistent, contradicting the definition of an inp-pattern. This proves the claim.

Note that, by the pigeonhole principle, we may assume that either (i) $\epsilon_\alpha, \epsilon'_\alpha \in r$ for all $\alpha < \kappa$, (ii) $\epsilon_\alpha \in r$, $\epsilon'_\alpha \in v_\alpha$ for all $\alpha < \kappa$, or (iii) $\epsilon_\alpha, \epsilon'_\alpha \in v_\alpha$ for all $\alpha < \kappa$. Case (i) is impossible by the choice of γ and, again by the pigeonhole principle, we may assume that if we are in case (ii), that ϵ_α is constant for all α. Then by rectification, we know that when $\alpha < \alpha'$, $\epsilon_\alpha \leq \epsilon_{\alpha'}$ and $\epsilon'_\alpha < \epsilon'_{\alpha'}$. Because for all $\alpha < \kappa$, the $c_{\alpha,k}$ are pairwise distinct and k varies, the set of formulas

$$\{(f_{\epsilon_\alpha \epsilon'_\alpha} \circ p_{\epsilon'_\alpha})(x) = c_{\alpha,k} : k < \omega\}$$

is 2-inconsistent. Moreover, if $g : \kappa \to \omega$ is a function, the partial type

$$\{(f_{\epsilon_\alpha \epsilon_{\alpha'}} \circ p_{\epsilon'_\alpha})(x) = c_{\alpha,g(\alpha)} : \alpha < \kappa\}$$

is implied by $\{\varphi_\alpha(x; a_{\alpha,g(\alpha)}) : \alpha < \kappa\}$ and is therefore consistent. It follows that $((f_{\epsilon_\alpha \epsilon'_\alpha} \circ p_{\epsilon'_\alpha})(x) = y_\alpha)_{\alpha<\kappa}, (c_{\alpha,k})_{\alpha<\kappa,k<\omega}$ is an inp-pattern. By Lemma 2.4.4, $f(\{\epsilon'_\alpha, \epsilon'_{\alpha'}\}) = 1$ for all $\alpha < \alpha'$. Therefore $H = \{\epsilon'_\alpha : \alpha < \kappa\}$ is a homogeneous set for f. $\square$

Case 2: $\kappa_{\mathbf{sct}} = \kappa^+$

In this subsection, we show that if $\kappa_{sct}(T^*_{\kappa,f}) = \kappa^+$ then f satisfies a homogeneity property inconsistent with f being a strong coloring. In particular, we will show that if this homogeneity property fails, then for any putative sct-pattern of height κ, there are two incomparable elements in $\omega^{<\kappa}$ which index compatible formulas, contradicting the inconsistency condition in the definition of an sct-pattern. This step is accomplished by relating consistency of the relevant formulas to an amalgamation problem in finite structures. The following lemma describes the relevant amalgamation problem:

Lemma 2.4.7. *Suppose given:*

- *Finite sets $w, w' \subset \kappa$ with $w \cap w' = v$ so that for all $\alpha \in v$, $\beta \in w \setminus v$, $\gamma \in w' \setminus v$, we have $\alpha < \beta < \gamma$ and $f(\{\beta, \gamma\}) = 1$.*
- *Structures $A \in \mathbb{K}_{w\cup w'}$, $B = \langle d, A\rangle_{L_w} \in \mathbb{K}_w$, $C = \langle e, A\rangle_{L_{w'}} \in \mathbb{K}_{w'}$ so that the map sending $d \mapsto e$ induces an isomorphism of L_v-structures over A between $\langle d, A\rangle_{L_v}$ and $\langle e, A\rangle_{L_v}$.*

Then there is $D = \langle f, A\rangle^D_{L_{w\cup w'}} \in \mathbb{K}_{w\cup w'}$ extending A so that $l(f) = l(d) = l(e)$ and $\langle f, A\rangle^D_{L_w} \cong B$ over A and $\langle f, A\rangle^D_{L_{w'}} \cong C$ over A via the isomorphisms over A sending $f \mapsto d$ and $f \mapsto e$, respectively.

Proof. Let f be a tuple of formal elements with $l(f) = l(d)(= l(e))$ with L_w and $L_{w'}$ interpreted so that $\langle f, A\rangle_{L_w}$ extends A and is isomorphic over A to B, and so that $\langle f, A\rangle_{L_{w'}}$ extends A and is isomorphic over A to C. Let D have underlying set

$$\langle f, A\rangle_{L_w} \cup \langle f, A\rangle_{L_{w'}} \cup \{*_{\alpha,c} : \alpha \in w \setminus v, c \in P_\gamma(\langle f, A\rangle_{L_{w'}}) \setminus P_\gamma(A)\},$$

where γ is the least element of $w' \setminus v$. We must give D an $L_{w\cup w'}$-structure. Let the elements of A, $\langle f, A\rangle_{L_w}$, $\langle f, A\rangle_{L_{w'}}$ in D inherit the interpretations from these respective structures. Interpret the predicates on the new formal elements $*_{\alpha,c}$ by ensuring $P_\alpha(*_{\alpha,c})$ holds and no other predicates hold on this element for all $\alpha \in w$ and $c \in P_\gamma(\langle f, A\rangle_{L_{w'}}) \setminus P_\gamma(A)$. Given $\alpha \in w \setminus v$ and $c \in P_\gamma(\langle f, A\rangle_{L_{w'}}) \setminus P_\gamma(A)$, interpret $f^D_{\alpha\gamma}(c) = *_{\alpha,c}$ and for any $\beta \in w' \setminus v$, define $f^D_{\alpha\beta} = f^D_{\alpha\gamma} \circ f^D_{\gamma\beta}$ on P^D_β and the identity on $D \setminus P^D_\beta$. If $\alpha \in w \setminus v$ and $\xi \in v$, interpret $f^D_{\xi\alpha}$ so that $f^D_{\xi\alpha}(*_{\alpha,c}) = f^D_{\xi\gamma}(c)$. Finally, interpret each function of the form p_β for $\beta \in w \cup w'$ to be the identity on the $*_{\alpha,c}$. This completes the definition of the $L_{w\cup w'}$-structure on D.

Now we must check that $D \in \mathbb{K}_{w\cup w'}$. It is easy to check that axioms $(1)-(3)$ are satisfied in D. As $f(\{\alpha, \beta\}) = 1$ for all $\alpha \in w \setminus v, \beta \in w' \setminus v$, the only possible counterexample to axiom (4) can occur when $\xi \in v$, $\beta \in (w \cup w') \setminus v$ and $f(\{\xi, \beta\}) = 0$. As the formal elements $*_{\alpha,c}$ are not in the image of O under the p_α, it follows that a counterexample to axiom (4) must come from a counter-example either in B or C, which is impossible. So $D \in \mathbb{K}_{w\cup w'}$, which completes the proof. $\square$

Lemma 2.4.8. *Suppose $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$, $(a_\eta)_{\eta\in\omega^{<\kappa}}$ is a rectified sct-pattern such that $l(x)$ is minimal among sct-patterns of height κ. Then for all $\alpha < \kappa$, $\varphi_\alpha(x; y_\alpha) \vdash (x)_l \in O \cup \bigcup_{i<n} P_{\alpha_i}$ for all $l < l(x)$.*

Proof. It is easy to see that if, for some $l < l(x)$ and all $\alpha < \kappa$, $\varphi_\alpha(x; y_\alpha) \vdash (x)_l \notin O \cup \bigcup_{i<n} P_{\alpha_i} \cup \bigcup_{j<m} P_{\beta_{\alpha,j}}$, then the only relations that $\varphi_\alpha(x; y_\alpha)$ can assert between $(x)_l$ and the elements of y_α and the other elements of x are equalities and inequalities allowing us to find an sct-pattern in fewer variables, contradicting minimality (or if $l(x) = 1$ a contradiction). So there is some $\alpha < \kappa$ and $j < m$ so that $\varphi_\alpha(x; y_\alpha) \vdash (x)_l \in P_{\beta_{\alpha,j}}$ and therefore, for all $\alpha' \neq \alpha$, $\varphi_{\alpha'}(x; y_{\alpha'})$ implies that $(x)_l$ is not in any of the unary predicates of $L_{w_{\alpha'}}$, as $\beta_{\alpha,j}$ is outside the root of the Δ-system. So $(\varphi_{\alpha'}(x; y_{\alpha'}) : \alpha' < \kappa, \alpha' \neq \alpha)$ forms an sct-pattern which falls into the first case considered, a contradiction. $\square$

Proposition 2.4.9. *Suppose $(\varphi_\alpha(x; y_\alpha) : \alpha < \kappa)$ is a rectified sct-pattern such that $l(x)$ is minimal among sct-patterns of height κ and whose associated Δ-system is $(w_\alpha : \alpha < \kappa)$, with $v_\alpha = w_\alpha \setminus r$, where r is the root. Then there is γ so that for any α, α' with $\gamma < \alpha < \alpha' < \kappa$ there is $\xi \in v_\alpha, \zeta \in v_{\alpha'}$ so that $f(\{\xi, \zeta\}) = 0$.*

Proof. Suppose not. By Lemma 2.4.8, we know that up to a relabeling of the variables, there is a $k \leq l(x)$ so that $\varphi_\alpha(x; y_\alpha) \vdash (x)_l \in \bigcup_{i<n} P_{\alpha_i}$ for $l < k$ and $\varphi_\alpha(x; y_\alpha) \vdash (x)_l \in O$ for $l \geq k$. Choose $\gamma < \kappa$ so that if $\delta < \kappa$ and $\varphi_\delta(x; a_{0^\delta}) \vdash f_{\alpha_i\alpha_{i'}}((x)_l) = c$ for $l < k$, or $\varphi_\delta(x; a_{0^\delta}) \vdash p_{\alpha_i}((x)_l) = c$ or $\varphi_\delta(x; a_{0^\delta}) \vdash (f_{\alpha_i\alpha_{i'}} \circ p_{\alpha_{i'}})((x)_l) = c$ for $l \geq k$, for some c, then this is implied by $\varphi_{\delta'}(x; 0^{\delta'})$ for some $\delta' < \gamma$ (possible as the root is finite). By assumption, there are α, α' with $\gamma < \alpha < \alpha' < \kappa$ so that $f(\{\xi, \zeta\}) = 1$ for all $\xi \in v_\alpha, \zeta \in v_{\alpha'}$. Choose $\eta \in \omega^\alpha$, $\nu \in \omega^{\alpha'}$ so that $\eta \perp \nu$. Let $A = \langle a_\eta, a_\nu\rangle_{L_{w_\alpha\cup w_{\alpha'}}}$ be the finite $L_{w_\alpha\cup w_{\alpha'}}$-structure generated by a_η and a_ν. By assumption and quantifier-elimination, it is possible to choose d with $\varphi_\alpha(d; a_\eta)$ and e with $\varphi_{\alpha'}(e; a_\nu)$ so that $\text{tp}_{L_r}(d/A) = \text{tp}_{L_r}(e/A)$. Let $B = \langle d, A\rangle_{L_{w_\alpha}}$ and $C = \langle e, A\rangle_{L_{w_{\alpha'}}}$. By Lemma 2.4.7, there is a $D \in \mathbb{K}_{w_\alpha\cup w_{\alpha'}}$ so that $D = \langle g, A\rangle^D_{L_{w_\alpha\cup w_{\alpha'}}}$ so that $l(g) = l(d) = l(e)$ and $\langle g, A\rangle_{L_{w_\alpha}} \cong B$ over A and $\langle g, A\rangle_{L_{w_{\alpha'}}} \cong C$ over A. It follows by model-completeness that in $\mathbb{M}$, $g \models \{\varphi_\alpha(x; a_\eta), \varphi_{\alpha'}(x; a_\nu)\}$, contradicting the definition of sct-pattern. This completes the proof. $\square$

Conclusion

Theorem 2.4.10. *There is a theory T so that $\kappa_{cdt}(T) \neq \kappa_{sct}(T) + \kappa_{inp}(T)$. Moreover, it is consistent with ZFC that for every regular uncountable κ, there is a theory T with $|T| = \kappa$ and $\kappa_{cdt}(T) \neq \kappa_{sct}(T) + \kappa_{inp}(T)$.*

Proof. If κ is regular and uncountable satisfying $\mathrm{Pr}_1(\kappa, \kappa, 2, \aleph_0)$, then choose $f : [\kappa]^2 \to 2$ witnessing $\mathrm{Pr}_1(\kappa, \kappa, 2, \aleph_0)$. There can be no homogeneous set of size κ for f, using Observation 2.2.9, and, moreover, there can be no collection $(v_\alpha : \alpha < \kappa)$ of disjoint finite sets so that, given $\alpha < \alpha' < \kappa$, there are $\xi \in v_\alpha, \zeta \in v_{\alpha'}$ so that $f(\{\xi, \zeta\}) = 0$. Let $T = T^*_{\kappa,f}$. Then $\kappa_{cdt}(T) = \kappa^+$, by Proposition 2.4.3, but $\kappa_{\mathrm{sct}}(T) < \kappa^+$ and $\kappa_{\mathrm{inp}}(T) < \kappa^+$ by Proposition 2.4.9 and Proposition 2.4.6 respectively. By Fact 2.2.11 and Observation 2.2.9, $\mathrm{Pr}_1(\lambda^{++}, \lambda^{++}, 2, \aleph_0)$ holds for any regular uncountable λ. Then $T = T^*_{\kappa,f}$ gives the desired theory, for $\kappa = \lambda^{++}$ and any f witnessing $\mathrm{Pr}_1(\lambda^{++}, \lambda^{++}, 2, \aleph_0)$. For the "moreover" clause, note that ZFC is equiconsistent with ZFC + GCH + "there are no inaccessible cardinals" which entails that every regular uncountable cardinal is a successor. By Fact 2.2.10 this implies that $\mathrm{Pr}_1(\kappa, \kappa, 2, \aleph_0)$ holds for all regular uncountable cardinals κ, which completes the proof. $\square$

Remark 2.4.11. In Chapter 1, it was proved that $\kappa_{\mathrm{cdt}}(T) = \kappa_{\mathrm{inp}}(T) + \kappa_{\mathrm{sct}}(T)$ for any countable theory T. The above theorem shows that in a certain sense, this result is best possible.

2.5 Compactness of ultrapowers

We say an ultrafilter $\mathcal{D}$ on I is *regular* if there is a collection of sets $\{X_\alpha : \alpha < |I|\} \subset \mathcal{D}$ such that for all $t \in I$, the set $\{\alpha : t \in X_\alpha\}$ is finite and $\mathcal{D}$ is *uniform* if all sets in $\mathcal{D}$ have cardinality $|I|$. In this section we study the decay of saturation in regular ultrapowers. Given a theory T, we start with a regular uniform ultrafilter $\mathcal{D}$ on λ and a λ^{++}-saturated model $M \models T$. We then consider whether the ultrapower $M^\lambda/\mathcal{D}$ is λ^{++}-compact. Shelah has shown [She90, VI.4.7] that if T is not simple, then in this situation $M^\lambda/\mathcal{D}$ will not be λ^{++}-compact and asked whether an analogous result holds for theories T with $\kappa_{\mathrm{inp}}(T) > \lambda^+$. We will show by direct construction that $\kappa_{\mathrm{inp}}(T) > \lambda^+$ does not suffice but, by modifying an argument due to Malliaris and Shelah [MS15b, Claim 7.5], $\kappa_{\mathrm{sct}}(T) > \lambda^+$ is sufficient to obtain a decay in compactness, by levaraging the finite square principles of Kennedy and Shelah [KS02].

A counterexample

Fix κ a regular uncountable cardinal. Let $L'_\kappa = \langle O, P_\alpha, p_\alpha : \alpha < \kappa \rangle$ be a language where O and each P_α is a unary predicate and each p_α is a unary function. Define a theory T'_κ to be the universal theory with the following as axioms:

1. O and the P_α are pairwise disjoint.
2. p_α is a function so that $(\forall x \in O)[p_\alpha(x) \in P_\alpha]$ and $(\forall x \notin O)[p_\alpha(x) = x]$.

Given a finite set $w \subset \kappa$, define $L'_w = \langle O, P_\alpha, p_\alpha : \alpha \in w \rangle$. Let $\mathbb{K}'_w$ denote the class of finite models of $T'_\kappa \restriction L'_w$.

Lemma 2.5.1. *Suppose $w \subset \kappa$ is finite. Then $\mathbb{K}'_w$ is a Fraïssé class*

Proof. The axioms of $T'_\kappa \restriction L_w$ are universal so HP is clear. As we allow the empty structure to be a model, JEP follows from AP. For AP, we reduce to the case where $A, B, C \in \mathbb{K}'_w$, A is a substructure of both B and C and $B \cap C = A$. Define an L'_w-structure D on $B \cup C$ by taking unions of the relations and functions as interpreted on B and C. It is easy to see that $D \in \mathbb{K}'_w$, so we're done. □

By Fraïssé theory, for each finite $w \subset \kappa$, there is a unique countable ultrahomogeneous L'_w-structure with age $\mathbb{K}'_w$. Let $T^\dagger_w$ denote its theory.

Lemma 2.5.2. *Suppose v and w are finite sets with $w \subset v \subset \kappa$. Then $T^\dagger_w \subset T^\dagger_v$.*

Proof. By induction, it suffices to consider the case when $v = w \cup \{\gamma\}$ for some $\gamma \in \kappa \setminus w$. By Fact 2.2.7, we must show (1) that $A \in \mathbb{K}'_v$ if and only if there is $D \in \mathbb{K}'_w$ so that A is an L'_v-substructure of $D \restriction L_v$ and (2) that whenever $A, B \in \mathbb{K}'_w$, $\pi : A \to B$ is an embedding, and $C \in \mathbb{K}'_v$ satisfies $C = \langle A \rangle^C_{L'_v}$ then there is $D \in \mathbb{K}'_v$ so that $D = \langle B \rangle^D_{L'_v}$ and π extends to an L'_v-embedding $\pi : C \to D$.

For (1), it is clear from definitions that if $D \in \mathbb{K}_w$ then $D \restriction L_v \in \mathbb{K}_v$. Given $A \in \mathbb{K}_v$, we may construct a suitable L_w-structure D as follows: let the underlying set of D be $A \cup \{*\}$ and interpret the predicates and functions to extend their interpretations on A and so that $P^D_\gamma = \{*\}$ and p^D_γ is the identity on the complement of O^D $(= O^A)$ and the constant function with value $*$ on O^D. Clearly $D \in \mathbb{K}_w$ and A is an L_v-substructure of $D \restriction L_v$.

For (2), suppose $A, B \in \mathbb{K}'_w$, $\pi : A \to B$ is an embedding, and $C \in \mathbb{K}'_v$ satisfies $C = \langle A \rangle^C_{L'_v}$. The requirement that $C = \langle A \rangle^C_{L'_v}$ entails that any points of $C \setminus A$ lie in P_γ. Write $O(B) = \pi(O(A)) \sqcup E$. Define an L'_v-structure D whose underlying set is $B \cup P_\gamma(A) \cup \{*_e : e \in E\}$. Interpret the relations on D so that all symbols of L'_w agree with their interpretations on B and define $P_\gamma(D) = P_\gamma(A) \cup \{*_e : e \in E\}$. Finally, define p^D_γ by

$$p^D_\gamma(x) = \begin{cases} p^C_\gamma(a) & \text{if } x = \pi(a) \\ *_x & \text{if } x \notin \pi(O(C)). \end{cases}$$

Clearly $D \in \mathbb{K}'_v$. Extend π to a map $\pi : C \to D$ by defining π to be the identity on $P_\gamma(C)$. We claim this is an embedding: note that for all $x \in O(C)$, $p^D_\gamma(\pi(x)) = p^C_\gamma(x) = \pi(p^C_\gamma(x))$ and π obviously respects all other structure from L'_w. $\square$

Define the theory $T^\dagger_\kappa$ to be the union of $T^\dagger_w$ for all finite $w \subset \kappa$. This is a complete stable L'_κ-theory with quantifier-elimination, as these properties are inherited from the T^*_w. Fix a monster $\mathbb{M} \models T^\dagger_\kappa$ and work there.

Proposition 2.5.3. $\kappa_{inp}(T^\dagger_\kappa) = \kappa^+$.

Proof. For each $\alpha < \kappa$, choose distinct $a_{\alpha,\beta} \in P_\alpha(\mathbb{M})$ for all $\beta < \omega$. It is easy to check that, for all functions $f : \kappa \to \omega$, $\{p_\alpha(x) = a_{\alpha,f(\alpha)} : \alpha < \kappa\}$ is consistent and, for all $\alpha < \kappa$, $\{p_\alpha(x) = a_{\alpha,\beta} : \beta < \omega\}$ is 2-inconsistent. Thus $(p_\alpha(x) = y_\alpha : \alpha < \kappa)$, $(a_{\alpha,\beta})_{\alpha<\kappa,\beta<\omega}$ forms an inp-pattern of height κ so $\kappa_{\text{inp}}(T^\dagger_\kappa) \geq \kappa^+$. The upper bound $\kappa_{\text{inp}}(T^\dagger_\kappa) \leq \kappa^+$ follows from the stability of $T^\dagger_\kappa$. $\square$

Proposition 2.5.4. *Suppose $\mathcal{D}$ is a regular ultrafilter on λ, $\kappa = \lambda^+$, and $M \models T^\dagger_\kappa$ is λ^{++}-saturated. Then $M^\lambda/\mathcal{D}$ is λ^{++}-saturated.*

Proof. Suppose $A \subseteq M^\lambda/\mathcal{D}$, $|A| = \kappa = \lambda^+$. To show that any $q(x) \in S^1(A)$ is realized, we have three cases to consider:

1. $q(x) \vdash x \in P_\alpha$ for some $\alpha < \kappa$
2. $q(x) \vdash x \notin O$ and $q(x) \vdash x \notin P_\alpha$ for all $\alpha < \kappa$
3. $q(x) \vdash x \in O$.

It suffices to consider q non-algebraic and $A = \text{dcl}(A)$. In case (1), $q(x)$ is implied by $\{P_\alpha(x)\} \cup \{x \neq a : a \in A\}$ and in case (2), $q(x)$ is implied by $\{\neg O(x) \wedge \neg P_\alpha(x) : \alpha < \kappa\} \cup \{x \neq a : a \in A\}$. To realize $q(x)$ in case (1), for each $t \in \lambda$, choose some $b_t \in P_\alpha(M)$ such that $b_t \neq a[t]$ for all $a \in A$, which is possible by the λ^{++}-saturation of M and the fact that $|A| = \lambda^+$. Let $b = \langle b_t \rangle_{t\in\lambda}/\mathcal{D}$. By Los's theorem, $b \models q$. Realizing q in case (2) is entirely similar.

So now we show how to handle case (3). Fix some complete type $q(x) \in S_1(A)$ such that $q(x) \vdash x \in O$. First, we note that by possibly growing A by κ many elements, we may assume that

$$q(x) = \{x \in O\} \cup \{x \neq a : a \in O(A)\} \cup \{p_\alpha(x) = c_\alpha\},$$

since, for each $\alpha < \kappa$, either $q(x) \vdash p_\alpha(x) = c_\alpha$ for some c_α, or it only proves inequations of this form. In the latter case, we can choose some element $c_\alpha \in P_\alpha(M^\lambda/\mathcal{D})$ not in A (possible by case (1) above) and extend $q(x)$ by adding the formula $p_\alpha(x) = c_\alpha$, which will then imply all inequations of the form $p_\alpha(x) \neq a$ for any $a \in A$, and this clearly remains finitely satisfiable. So now given q in the form described above, let $X_t = \{\alpha < \kappa : M \models P_\alpha(c_\alpha[t])\}$. Let the type $q_t(x)$ be defined by

$$q_t(x) = \{x \in O\} \cup \{x \neq a[t] : a \in O(A)\} \cup \{p_\alpha(x) = c_\alpha[t] : \alpha \in X_t\}.$$

By construction, if $\alpha \neq \alpha' \in X_t$ then $M \models P_\alpha(c_\alpha[t]) \wedge P_{\alpha'}(c_{\alpha'}[t])$ so this type is consistent and over a parameter set of size at most κ, hence realized by some $b_t \in M$. Let $b = \langle b_t \rangle_{t\in\lambda}/\mathcal{D}$ and let J_α be defined by $J_\alpha = \{t \in \lambda : M \models P_\alpha(c_\alpha[t])\}$. As $q(x)$ is a consistent type, $J_\alpha \in \mathcal{D}$ and, by construction, $J_\alpha \subseteq \{t \in \lambda : M \models p_\alpha(b_t) = c_\alpha[t]\}$ so $M^\lambda/\mathcal{D} \models p_\alpha(b) = c_\alpha$. It is obvious that b satisfies all of the other formulas of q so we're done. □

Corollary 2.5.5. *Suppose T is a complete theory, $|I| = \lambda$, $\mathcal{D}$ on I is a regular ultrafilter, and $M \models T$ is a λ^{++}-saturated model of T. The condition that $\kappa_{inp}(T) > |I|^+$ is, in general, not sufficient to guarantee that $M^I/\mathcal{D}$ is not λ^{++}-compact. In particular, the condition that $\kappa_{cdt}(T) > |I|^+$ is, in general, not sufficient to guarantee that $M^I/\mathcal{D}$ is not λ^{++}-compact.*

Proof. Given λ, I with $|I| = \lambda$, and $\mathcal{D}$, a regular ultrafilter on I, choose any λ^{++}-saturated model of $T^\dagger_{\lambda^+}$. By Lemma 2.5.3, $\kappa_{\mathrm{cdt}}(T^\dagger_{\lambda^+}) \geq \kappa_{\mathrm{inp}}(T^\dagger_{\lambda^+}) = \lambda^{++} > |I|^+$, but, by Proposition 2.5.4, $M^I/\mathcal{D}$ is λ^{++}-saturated and hence λ^{++}-compact. □

Loss of saturation from large sct-patterns

If T is not simple, then it has either the tree property of the first kind or the second kind - Shelah argues in [She90, VI.4.7] by demonstrating that either property results in a decay of saturation with an argument tailored to each property. The preceding section demonstrates that the analogy between TP_2 and $\kappa_{\mathrm{inp}}(T) > |I|^+$ breaks down, but we show that the analogy between TP_1 and $\kappa_{\mathrm{sct}}(T) > |I|^+$ survives. The following is a straightforward adaptation of the argument of [MS15b, Claim 8.5]:

Fact 2.5.6. [KS02, Lemma 4] Suppose $\mathcal{D}$ is a regular uniform ultrafilter on λ and $\lambda = \lambda^{<\lambda}$. There is an array of sets $\langle u_{t,\alpha} : t < \lambda, \alpha < \lambda^+\rangle$ satisfying the following properties:

1. $u_{t,\alpha} \subseteq \alpha$
2. $|u_{t,\alpha}| < \lambda$
3. $\alpha \in u_{t,\beta} \implies u_{t,\beta} \cap \alpha = u_{t,\alpha}$
4. if $u \subseteq \lambda^+$, $|u| < \aleph_0$ then $\{t < \lambda : (\exists \alpha)(u \subseteq u_{t,\alpha})\} \in \mathcal{D}$.

Theorem 2.5.7. *Suppose $|I| = \lambda$ and $\lambda = \lambda^{<\lambda}$. Suppose $\kappa_{sct}(T) > |I|^+$, M is an $|I|^{++}$-saturated model of T and $\mathcal{D}$ is a regular ultrafilter over I. Then $M^I/\mathcal{D}$ is not $|I|^{++}$-compact.*

Proof. Let $(\varphi_\alpha(x; y_\alpha) : \alpha < \lambda^+)$, $(a_\eta)_{\eta\in\lambda^{<\lambda^+}}$ be an sct-pattern. We may assume $l(y_\alpha) = k$ for all $\alpha < \lambda^+$. Let $\langle u_{t,\epsilon} : t < \lambda, \alpha < \lambda^+\rangle$ be given as by Fact 2.5.6. By induction on $\alpha < \lambda^+$, we'll construct $\langle \eta_{t,\alpha} : \alpha < \lambda^+\rangle$ such that $\eta_{t,\alpha} \in \lambda^\alpha$ and $\eta_{t,\alpha} \trianglelefteq \eta_{t,\beta} \iff \alpha \in u_{t,\beta}$: suppose $\langle \eta_{t,\beta} : \beta < \alpha\rangle$ has been constructed. The set $\{\eta_{t,\beta} : \beta \in u_{t,\alpha}\}$ is contained in a path since, if $\beta < \beta'$ are elements of $u_{t,\alpha}$ then $\beta \in u_{t,\alpha} \cap \beta' = u_{t,\beta'}$ so $\eta_{t,\beta} \trianglelefteq \eta_{t,\beta'}$ by induction. Then we can pick $\eta_{t,\alpha} \in \lambda^\alpha$ so that $\eta_{t,\beta} \trianglelefteq \eta_{t,\alpha}$ if and only if $\beta \in u_{t,\alpha}$. For each $\alpha < \lambda^+$ we thus have an element $c_\alpha \in M^\lambda/\mathcal{D}$ given by $c_\alpha = \langle c_\alpha[t] : t < \lambda\rangle/\mathcal{D}$ where $c_\alpha[t] = a_{\eta_{t,\alpha}} \in M$.

Claim: $p(x) := \{\varphi_\alpha(x; c_\alpha) : \alpha < \lambda^+\}$ is consistent.

Fix any finite $u \subseteq \lambda^+$. If for some $t < \lambda$ and $\alpha < \lambda^+$, we have $u \subseteq u_{t,\alpha}$ then $\{\eta_{t,\beta} : \beta \in u\} \subseteq \{\eta_{t,\beta} : \beta \in u_{t,\alpha}\}$ which is contained in a path, hence $\{\varphi_\beta(x; c_\beta[t]) : \beta \in u\} = \{\varphi_\beta(x; a_{\eta_{t,\beta}}) : \beta \in u\}$ is consistent by definition of an sct-pattern. We know $\{t < \lambda : (\exists\alpha)(u \subseteq u_{t,\alpha})\} \in \mathcal{D}$ so the claim follows by Los's theorem and compactness.

Suppose $b = \langle b[t]\rangle_{t\in\lambda}/\mathcal{D}$ is a realization of p in $M^\lambda/\mathcal{D}$. For each $\alpha < \lambda^+$ define $J_\alpha = \{t < \lambda : M \models \varphi_\alpha(b[t], c_\alpha[t])\} \in \mathcal{D}$. For each α, pick $t_\alpha \in J_\alpha$. The map $\alpha \mapsto t_\alpha$ is regressive on the stationary set of α with $\lambda \leq \alpha < \lambda^+$. By Fodor's lemma, there's some t_* so that the set $S = \{\alpha < \lambda^+ : t_\alpha = t_*\}$ is stationary. Therefore $p_*(x) = \{\varphi_\alpha(x; a_{\eta_{t_*,\alpha}}) : \alpha \in S\}$ is a consistent partial type in M so $\{\eta_{t_*,\alpha} : \alpha \in S\}$ is contained in a path, by definition of sct-pattern. Choose an $\alpha \in S$ so that $|S \cap \alpha| = \lambda$. Then, by choice of the $\eta_{t,\alpha}$, we have $\beta \in S \cap \alpha$ implies $\eta_{t_*,\beta} \trianglelefteq \eta_{t_*,\alpha}$ and therefore $\beta \in u_{t_*,\alpha}$. This shows $|u_{t_*,\alpha}| \geq \lambda$, a contradiction. $\square$

Part III

$NSOP_1$ in detail

Chapter 3

Kim-Independence

This chapter is joint work with Itay Kaplan.

3.1 Introduction

The class of simple theories was one of the first classes of unstable theories to receive extensive study. The starting point is *Classification Theory*, where, in the course of studying stable theories, Shelah isolates *local character* as a key property of non-forking independence and observes a dichotomy in the way local character can fail, a theorem we now recognize as saying that a non-simple theory must have the tree property of the first or second kind [She90, Theorem III.7.11]. Shortly after the publication of the first edition of [She90], Shelah defined the class of simple theories and characterized them in terms of a certain chain condition of the Boolean algebra of non-weakly dividing formulas, which in turn led to consistency results on their saturation spectra [She80]. The aim of that work was to obtain an 'outside' set-theoretic definition of the class to support the claim that simplicity marked a dividing line. In separate developments, questions concerning concrete examples created the need for new methods to treat unstable structures. Hrushovski and Pillay used local stability and S_1-rank in the study of the definability of groups in pseudo-finite and PAC fields in [HP94], and these methods were situated in the broader context of PAC structures studied by Hrushovski [Hru91], where an independence theorem was proved. Moreover, Lachlan's far-reaching theory of smoothly approximated structures furnished examples of tame unstable theories. After Kantor, Liebeck, and Macpherson [KLM89] classified the primitive smoothly approximable structures, Cherlin and Hrushovski [CH03] used stability theoretic methods concerning independence and amalgamation to describe how these primitive pieces fit together to form a quasi-finite structure.

Kim's thesis and subsequent work by Kim and Pillay showed how to regard these developments as instances of a common theory, with non-forking independence at its center [Kim98], [KP97]. Kim proved that in a simple theory, forking and dividing coincide, non-forking independence is symmetric and transitive, and Kim and Pillay proved that the independence theorem holds over models. Moreover, Kim showed that symmetry and transitivity of non-forking both individually *characterize* the simple theories, and Kim and Pillay showed that any independence relation satisfying the basic properties of non-forking independence must actually coincide with non-forking independence, giving both a striking characterization of the simple theories and a powerful method for showing that a particular theory is simple, namely by observing that it has an independence relation of the right kind.

Here, we study the class of NSOP_1 theories. These are the theories which do not have the property SOP_1, which form a class of theories that properly contain the simple theories and which are contained inside the class of theories without the tree property of the first kind. SOP_1 was defined by Džamonja and Shelah in their study of the $\trianglelefteq^*$-order [DS04] and later studied by Shelah and Usvyatsov in [SU08]. The NSOP_1 theories were characterized as the theories satisfying a weak independence theorem for invariant types in Chapter 1. This characterization provided a point of contact between the combinatorics of model-theoretic tree properties and the study of definability in particular algebraic examples. Chatzidakis [Cha99], [Cha02] studied independence in ω-free PAC fields and, more generally, Frobenius fields and showed that the independence theorem holds for these structures even though they are not simple. Similarly, Granger showed in his thesis that the model companion of the theory of infinite-dimensional vector spaces with a bilinear form is not simple but nonetheless comes equipped with a good notion of independence. The amalgamation criterion of Chapter 1 established that these structures have NSOP_1 theory by appealing to the existence of these independence relations, but what was missing was a theory of independence in NSOP_1 theories more generally. The purpose of this paper is to establish exactly such a theory.

One central tool in the study of forking in simple theories is Kim's lemma: in a simple theory, a formula divides over a set A if and only if it divides with respect to some Morley sequence over A if and only if it divides for all Morley sequences over A. In [CK12], this was shown to hold over models in NTP_2 theories, provided that the Morley sequence is a strict invariant Morley sequence. In the setting of NSOP_1 theories, we find a new phenomenon: forking which is *never* witnessed by a generic sequence. In fact, we show that any NSOP_1 theory with a universal witness to dividing must be simple (Proposition 3.8.7 below) and that forking need not equal dividing in an NSOP_1 theory. Nonetheless, we find that, by restricting attention to the forking that is witnessed by a generic sequence, one can recover many of the properties of forking in simple theories. We show moreover that this kind of simplicity at a generic scale is characteristic of NSOP_1 theories.

There is considerable freedom in the choice of notion of generic sequence. One suggestion which inspired our work is due to Kim, who proposed in his 2009 talk on NTP_1 theories [Kim09] that one might develop an independence theory for NTP_1 theories or a subclass therein by considering only formulas which divide with respect to every non-forking Morley sequence. Compared to invariance or finite satisfiability, forking is a relatively weak notion of independence and this notion proved unwieldy at the beginning stages of developing the theory presented here. However, Hrushovski's study of q-dividing [Hru12] and Malliaris and Shelah's characterization of NTP_1 theories in terms of higher formulas [MS15a] provided evidence that one might be able to build a theory around an investigation of formulas that divide with respect to a Morley sequence in a global invariant or finitely satisfiable type. Building off this work, we introduce the notion of Kim-dividing – a formula *Kim-divides* over a set A if it divides with respect to a Morley sequence in a global A-invariant type – and the associated notion of independence, *Kim-independence.* Our first observation is that a theory is NSOP_1 if and only if Kim-dividing satisfies a version of Kim's lemma over models, where a formula divides with respect to a Morley sequence in *some* global invariant type extending the type of the parameters if and only if it divides with respect to *every* Morley sequence in an appropriate invariant type.

From Kim's lemma for Kim-dividing, many familiar properties of non-forking independence follow: Kim-forking equals Kim-dividing, Kim-independence satisfies extension and a version of the chain condition, etc. In subsequent sections, we investigate additional properties of Kim-independence in NSOP_1 theories and prove that, in many cases, these properties are characteristic of NSOP_1. In Section 3.4 we observe a form of local character for Kim-independence in the context of NSOP_1 theories. In Section 3.5, we show additionally that Kim-independence is symmetric over models. The argument there centers upon the notion of a *tree Morley sequence* which is defined in terms of indiscernible trees. We show that tree Morley sequences always witness Kim-dividing and prove a version of the chain condition for them. In Section 3.6, we prove the independence theorem. In Section 3.7, we prove that in an NSOP_1 theory a formula Kim-divides over a model if and only if it divides with respect to every non-forking Morley sequence in the parameters and this too characterizes NSOP_1 theories. This means that Kim-independence could have been defined from the outset in essentially the way Kim proposed, but curiously, proving anything about this notion without making use of invariant types seems quite difficult. In Section 3.8, we state our main theorem: Kim's lemma for Kim-dividing, symmetry over models, and the independence theorem both hold in NSOP_1 theories and individually characterize NSOP_1 theories. We also show that the simple theories can be characterized in several new ways in terms of Kim-independence. In particular, we show that Kim-independence coincides with non-forking over models if and only if the theory is simple, which means that our theorems imply the corresponding facts for non-forking independence in a simple theory.

We conclude the paper with Section 3.9 where we describe Kim-independence explicitly in several concrete examples. We show it may be described in purely algebraic terms in the case of Frobenius fields, where Kim-independence turns out to coincide with *weak independence*, as defined by Chatzidakis. We also show that in Granger's two-sorted theory of a vector space over an algebraically closed field with a generic bilinear form, Kim-independence is closely related to Granger's Γ-independence and may be given a simple algebraic description. These results suggest the naturality and robustness of Kim-dividing, but also serve to explain the simplicity-like phenomena observed in these concrete examples on the basis of a general theory. We additionally describe a combinatorial example of a NSOP_1 theory, based on a variant of T^*_{feq} introduced by Džamonja and Shelah, which furnishes counter-examples to some *a priori* possible strengthenings of the results we prove. In particular, we give the first example of a simple non-cosimple type, answering a question of Chernikov [Che14], and the first example of an NSOP_3 theory in which every complete type has a global non-forking extension but forking does not equal dividing, answering a question of Conant [C+17].

3.2 Syntax

In this section we will define SOP_1 and prove its equivalence with a syntactic property of a different form. This will allow us to relate SOP_1 to dividing. We will often work with arrays and trees. Suppose $(c_{ij})_{i<\kappa,j<\lambda}$ is an array. Write $\bar{c}_i = (c_{i,j})_{j<\lambda}$ for the ith row of the array and $\bar{c}_{<i}$ for the sequence of rows with index less than i, i.e. $(\bar{c}_k)_{k<i}$. Suppose $\mathcal{T}$ is a tree, $(a_\eta)_{\eta\in\mathcal{T}}$ is a collection of tuples indexed by $\mathcal{T}$. We write $\trianglelefteq$ for the tree partial order and $<_{lex}$ for the lexicographic order on $\mathcal{T}$. For a node $\eta \in \mathcal{T}$, write $a_{\trianglelefteq\eta}$ for the sequence $\langle a_\nu : \nu \trianglelefteq \eta\rangle$, and likewise $a_{\triangleleft\eta}$ for $\langle a_\nu : \nu \triangleleft \eta\rangle$. We use the notation $a_{\trianglerighteq\eta}$ and $a_{\triangleright\eta}$ similarly. If the tree $\mathcal{T}$ is contained in $2^{<\kappa}$ or $\omega^{<\kappa}$, we write 0^α to denote the element of the tree of length α consisting of all zeros. Throughout the paper, T denotes a complete theory and $\mathbb{M} \models T$ is a monster model of T.

Definition 3.2.1. [DS04, Definition 2.2] The formula $\varphi(x; y)$ has SOP_1 if there is a collection of tuples $(a_\eta)_{\eta\in 2^{<\omega}}$ so that

- For all $\eta \in 2^\omega$, $\{\varphi(x; a_{\eta|\alpha}) : \alpha < \omega\}$ is consistent.
- For all $\eta \in 2^{<\omega}$, if $\nu \trianglerighteq \eta \frown \langle 0\rangle$, then $\{\varphi(x; a_\nu), \varphi(x; a_{\eta\frown 1})\}$ is inconsistent.

We say T is SOP_1 if some formula has SOP_1 modulo T. T is NSOP_1 otherwise.

The following lemma is close to Lemma I.5.2, but with a key strengthening which will allow us to relax the 2-inconstency in the definition of SOP_1 to a version with k-inconsistency.

Lemma 3.2.2. *Suppose $(c_{i,j})_{i<\omega,j<2}$ is an array where $c_{i,j} = (d_{ij}, e_{ij})$ for all i, j and $\chi_1(x; y)$ and $\chi_2(x; z)$ are formulas over C. Write $\psi(x; y, z)$ for $\chi_1(x; y) \wedge \chi_2(x; z)$ and suppose*

1. *For all $i < \omega$, $e_{i,0} \equiv_{Cc_{<i,0}e_{<i,1}} e_{i,1}$.*
2. *$\{\psi(x; c_{i,0}) : i < \omega\}$ is consistent.*
3. *$j \leq i \implies \{\chi_1(x; d_{i,0}), \chi_2(x; e_{j,1})\}$ is inconsistent.*

then T has SOP_1.

Proof. By adding constants, we may assume $C = \emptyset$. By Ramsey and compactness, we may assume $(\bar{c}_i)_{i<\omega}$ is a C-indiscernible sequence. By compactness again, we may extend the array to an array whose rows are indexed by the integers $(\bar{c}_i)_{i\in\mathbb{Z}}$. We will construct, for each $n < \omega$, a tree $(c_\eta)_{\eta\in 2^{\leq n}}$ so that

1. If $\nu \in 2^n$, then $\{\psi(x; c_{\nu|i}) : i \leq n\}$ is consistent.
2. If $\nu \in 2^{<n}$ and $\nu \frown \langle 0\rangle \triangleleft \eta\rangle$ then $\{\psi(x; c_\eta), \psi(x; c_{\nu\frown\langle 1\rangle})\}$ is inconsistent.

3. If $\eta \in 2^n$, $(c_\nu)_{\nu \trianglelefteq \eta} \equiv_{c_{<-n,0}e_{<-n,1}} (c_{i,0})_{-n \leq i \leq 0}$.

To define $(c_\eta)_{\eta \in 2^{\leq 0}}$, we put $c_\emptyset = c_{0,0}$. Now suppose we are given $S_n = (c_\eta)_{\eta \in 2^{\leq n}}$ satisfying the requirements. There is an automorphism σ taking $e_{-n,0}$ to $e_{-n,1}$ fixing $c_{<-n,0}e_{<-n,1}$. Define $S_{n+1} = (c'_\eta)_{\eta \in 2^{\leq n+1}}$ by $c'_\emptyset = c_{-(n+1),0}$ and, for all $\eta \in 2^{\leq n}$, $c'_{\langle 0 \rangle \frown \eta} = c_\eta$, $c'_{\langle 1 \rangle \frown \eta} = \sigma(c_\eta)$. Clearly all branches have the same type over $c_{<-(n+1),0}e_{<-(n+1),1}$ as $(c_{i,0})_{-(n+1) \leq i \leq 0}$. Write $c'_\eta = (d'_\eta, e'_\eta)$ for all $\eta \in 2^{\leq n+1}$. Now note that in both S_n and $\sigma(S_n)$ conditions (1) and (2) are preserved and that $\psi(x; c'_{\langle 1 \rangle})$ is inconsistent with $\psi(x; c'_{\langle 0 \rangle \frown \eta})$ for any $\eta \in 2^{\leq n}$ since $\chi_2(x; e_{-n,1}) \wedge \chi_1(x; d_\eta)$ is consistent if and only if $\chi_2(x; e_{-n,1}) \wedge \chi_1(x; d_{i,0})$ is consistent, for $i = l(\eta) - n$. Likewise, instantiating $\psi(x; y)$ along any branch through this tree yields something consistent: any branch in S_n or $\sigma(S_n)$ has the same type over $c_{-(n+1),0}$ as $(c_{i,0})_{-n \leq i \leq 0}$ and $\{\psi(x; c_{i,0}) : -(n+1) \leq i \leq 0\}$ is consistent. We conclude by compactness. $\square$

Lemma 3.2.3. *Suppose $\varphi(x; y)$ is a formula, k is a natural number, and $(\overline{c}_i)_{i \in I}$ is an infinite sequence with $\overline{c}_i = (c_{i,0}, c_{i,1})$ satisfying:*

1. *For all $i \in I$, $c_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}$.*
2. *$\{\varphi(x; c_{i,0}) : i \in I\}$ is consistent.*
3. *$\{\varphi(x; c_{i,1}) : i \in I\}$ is k-inconsistent.*

Then T has SOP_1.

Proof. By compactness and Ramsey, it suffices to prove this when $I = \mathbb{Q}$ – so suppose $(c_{i,0}, c_{i,1})_{i \in \mathbb{Q}}$ is an indiscernible sequence with $c_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}$, $\{\varphi(x; c_{i,0}) : i \in \mathbb{Q}\}$ is consistent, and $\{\varphi(x; c_{i,1}) : i \in \mathbb{Q}\}$ is k-inconsistent.

For integers $l < l'$, define a partial type $\Gamma_{l,l'}(x)$ by

$$\{\varphi(x; c_{i,0}) : i \in (l+m, l+m+1), m \in \omega, m < l' - l\} \cup \{\varphi(x; c_{l+m,1}) : m < l' - l, m \in \omega\}.$$

Let $\Gamma_{l,l}(x) = \emptyset$. Note that if $\Gamma_{l,l'}(x)$ is consistent then $\Gamma_{l+z,l'+z}(x)$ is consistent for any integer z by indiscernibility of the sequence $(\overline{c}_i)_{i \in \mathbb{Q}}$. Let $n \in \omega$ be maximal so that $\Gamma_{0,n}(x)$ is consistent. Note that $\Gamma_{0,0}(x)$ is consistent, as it is the empty partial type and we have

$$\Gamma_{0,k}(x) \vdash \{\varphi(x; c_{i,1}) : i \in \omega, i < k\},$$

which is inconsistent, so $0 \leq n < k$. So now we know $\Gamma_{-n,0}(x)$ is consistent and $\Gamma_{-n,1}(x) = \Gamma_{-n,0}(x) \cup \Gamma_{0,1}(x)$ is inconsistent. By indiscernibility and compactness, we may fix some integer $N > 0$ so that

$$\Gamma_{-n,0}(x) \cup \{\varphi(x; c_{0,1})\} \cup \{\varphi(x; c_{\frac{j+1}{N},0}) : j \in \omega, j < N-1\}$$

is inconsistent. Now choose $\Delta(x) \subseteq \Gamma_{-n,0}(x)$ finite so that

$$\Delta(x) \cup \{\varphi(x; c_{0,1})\} \cup \{\varphi(x; c_{\frac{j+1}{N},0}) : j \in \omega, j < N-1\}$$

is inconsistent. Let z indicate the tuple of variables $(y_0, \ldots, y_{N-2})$ and let $\chi(x; z)$ be the formula $\chi(x; z) = \bigwedge_{i<N} \varphi(x; y_i) \wedge \bigwedge \Delta(x)$. Let $(a_{i,j})_{i<\omega, j<2}$ be defined as follows:

$$a_{i,0} = (c_{i,0}; d_{i,0}) = (c_{i,0}; c_{i+\frac{1}{N},0}, \ldots, c_{i+\frac{N-1}{N},0}).$$

Now choose $d_{i,1}$ so that $c_{i,0}d_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}d_{i,1}$ – this is possible as $c_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}$. Then we put $a_{i,1} = (c_{i,1}, d_{i,1})$. Let $\psi(x; yz) = \varphi(x; y) \wedge \chi(x; z)$.

To conclude, we have to establish the following:

Claim: The array $(a_{i,j})_{i<\omega, j<2}$ and the formulas $\varphi(x; y)$, $\chi(x; z)$ satisfy the following:

1. $a_{i,0} \equiv_{a_{<i,0}, c_{<i,1}} a_{i,1}$.
2. $\{\psi(x; a_{i,0}) : i < \omega\}$ is consistent.
3. If $l \leq l'$ then $\{\varphi(x; c_{l,1}), \chi(x; d_{l',0})\}$ is inconsistent.

Proof of claim: (1) follows from the fact that $a_{i,0} \equiv_{\overline{c}_{<i}} a_{i,1}$ and both $a_{<i,0}$ and $c_{<i,1}$ are enumerated in $\overline{c}_{<i}$. Note that $\Gamma_{-n,0}(x)$ is consistent so, by indiscernibility,

$$\Gamma_{-n,0}(x) \cup \{\varphi(x; c_{i,0}) : i \in [0, \infty) \cap \mathbb{Q}\}$$

is consistent, which establishes (2). Finally, if $l \leq l'$, then $\{\varphi(x; c_{l,1}), \chi(x; d_{l',0})\}$ implies

$$\{\varphi(x; c_{l,1})\} \cup \{\varphi(x; c_{l'+\frac{j+1}{N}}) : j \in \omega, j < N-1\} \cup \Delta(x).$$

By indiscernibility of $(\overline{c}_i)_{i\in\mathbb{Q}}$ and the fact that $l \leq l'$, this set is consistent if and only if

$$\{\varphi(x; c_{0,1})\} \cup \{\varphi(x; c_{\frac{j+1}{N},0}) : j \in \omega, j < N-1\} \cup \Delta(x)$$

is consistent. As this latter set is inconsistent, this shows (3), which proves the claim. The lemma now follows by Lemma 3.2.2. □

Finally, we note that the criterion for SOP_1 from Lemma 4.2.2 is an equivalence. This was implicit in Chapter 1, at least in its 2-inconsistent version, but we think that the property described by Lemma 4.2.2 is, in most cases, the more fruitful way of thinking about SOP_1 and therefore worth making explicit.

Proposition 3.2.4. *The following are equivalent, for a complete theory T:*

1. *T has SOP_1.*
2. *There is a formula φ and an array $(c_{i,j})_{i<\omega, j<2}$ so that:*
   a) *$c_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}$ for all $i < \omega$.*
   b) *$\{\varphi(x; c_{i,0}) : i < \omega\}$ is consistent.*
   c) *$\{\varphi(x; c_{i,1}) : i < \omega\}$ is 2-inconsistent.*

3. *There is a formula φ and an array $(c_{i,j})_{i<\omega,j<2}$ so that:*

   a) $c_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}$ *for all* $i < \omega$.

   b) $\{\varphi(x; c_{i,0}) : i < \omega\}$ *is consistent.*

   c) $\{\varphi(x; c_{i,1}) : i < \omega\}$ *is* k*-inconsistent for some* k.

Proof. (3) $\implies$ (1) is Lemma 4.2.2.

(1) $\implies$ (2). This follows from the proof of Proposition I.5.6.

(2) $\implies$ (3) is obvious. □

Remark 3.2.5. Though the configurations described in (2) and (3) are not obviously preserved by expansion, SOP$_1$ as defined in Definition 3.2.1 clearly is. It follows, then, that one can take $(\overline{c}_i)_{i<\omega}$ to be indiscernible with respect to some Skolemization in the language L^{Sk} of T and, moreover, obtain $c_{i,0} \equiv^{L^{Sk}}_{\overline{c}_{<i}} c_{i,1}$ for all $i < \omega$ (in fact, this is what the proof of Proposition I.5.6 directly shows).

3.3 Kim-dividing

Averages and Invariant Types

Definition 3.3.1. A global type $q \in S(\mathbb{M})$ is called *A-invariant* if $b \equiv_A b'$ implies $\varphi(x; b) \in q$ if and only if $\varphi(x; b') \in q$. A global type q is *invariant* if there is some small set A such that q is A-invariant. If $q(x)$ and $r(y)$ are A-invariant global types, then the type $(q \otimes r)(x, y)$ is defined to be $\text{tp}(a, b/\mathbb{M})$ for any $b \models r$ and $a \models q|_{\mathbb{M}b}$. We define $q^{\otimes n}(x_0, \ldots, x_{n-1})$ by induction: $q^{\otimes 1} = q$ and $q^{\otimes n+1} = q(x_n) \otimes q^{\otimes n}(x_0, \ldots, x_{n-1})$. When M is a model, write $a \mathop{\smile\!\!\!\!\!\!\;\,|}^{i}_{M} b$ to mean $\text{tp}(a/Mb)$ extends to a global M-invariant type.

Fact 3.3.2. [Sim15, Chapter 2] Given a global A-invariant type q and positive integer n, $q^{\otimes n}$ is a well-defined A-invariant global type. If $N \supset A$ is an $|A|^+$-saturated model and $p \in S(N)$ satisfies $\varphi(x; b) \in p \iff \varphi(x; b') \in p$ whenever $b, b' \in N$ and $b \equiv_A b'$, then p extends uniquely to a global A-invariant type.

Definition 3.3.3. Suppose q is an A-invariant global type and I is a linearly ordered set. By a *Morley sequence in q over A of order type I*, we mean a sequence $(b_\alpha)_{\alpha\in I}$ such that for each $\alpha \in I$, $b_\alpha \models q|_{Ab_{<\alpha}}$ where $b_{<\alpha} = (b_\beta)_{\beta<\alpha}$. Given a linear order I, we will write $q^{\otimes I} = q^{\otimes I}(x_\alpha : \alpha \in I)$ for the A-invariant global type so that if $\overline{b} \models q^{\otimes I}$ then $b_\alpha \models q|_{\mathbb{M}b_{<\alpha}}$ for all $\alpha \in I$.

The above definition of $q^{\otimes I}$ generalizes the finite tensor product $q^{\otimes n}$ – given any global A-invariant type q and linearly ordered set I, one may easily show that $q^{\otimes I}$ exists and is A-invariant, by Fact 3.3.2 and compactness.

Definition 3.3.4. Let $I \subseteq \mathbb{M}^n$ be a collection of tuples, $A \subseteq \mathbb{M}$ a set, and $\mathcal{D}$ an ultrafilter over I. We define the *average type of* $\mathcal{D}$ *over* A to be the type defined by

$$\mathrm{Av}(\mathcal{D}, A) = \{\varphi(x; a) : a \in A \text{ and } \{b \in I : \mathbb{M} \models \varphi(b; a)\} \in \mathcal{D}\}.$$

Fact 3.3.5. [She90, Lemma 4.1] Let $I \subseteq \mathbb{M}^n$ be a collection of tuples and $\mathcal{D}$ an ultrafilter on I.

1. For every set C, $\mathrm{Av}(\mathcal{D}, C)$ is a complete type over C.
2. The global type $\mathrm{Av}(\mathcal{D}, \mathbb{M})$ is I-invariant.
3. For any model $M \models T$, if $p \in S^n(M)$, there is some ultrafilter $\mathcal{E}$ on M^n so that $p = \mathrm{Av}(\mathcal{E}, M)$.

One important consequence of Fact 4.2.4 for us is that every type over a model M extends to a global M-invariant type: given $p \in S(M)$, one chooses an ultrafilter $\mathcal{D}$ so that $\mathrm{Av}(\mathcal{D}, M) = p$. Then $\mathrm{Av}(\mathcal{D}, \mathbb{M})$ is a global type extending p which is M-invariant. In the arguments below, it will often be convenient to produce global invariant types through a particular choice of ultrafilter.

Fact 3.3.6. [CK12, Remark 2.16] Write $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^u_A b$ to mean that $\mathrm{tp}(a/Ab)$ is finitely satisfiable in A – the u is for "ultrafilter" as this is equivalent to asserting $\mathrm{tp}(a/Ab) = \mathrm{Av}(\mathcal{D}, Ab)$ for some ultrafilter $\mathcal{D}$ on A. The relation $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^u$ satisfies both left and right extension over models:

1. (Left extension) If M is a model and $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^u_M b$ then for all d, there is some $b' \equiv_{Ma} b$ so that $ad \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^u_M b'$.
2. (Right extension) If M is a model and $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^u_M b$ then for all c, there is some $a' \equiv_{Mb} a$ so that $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^u_M bc$.

Definition 3.3.7. Suppose $M \models T$ and $\overline{a} = (a_i)_{i<\omega}$ is an M-indiscernible sequence. A global M-invariant type $q \supseteq \mathrm{tp}(\overline{a}/M)$ is called an *indiscernible type* if whenenver $\overline{a}' \models q$, $\overline{a}'$ is $\mathbb{M}$-indiscernible.

Definition 3.3.8. A collection of sequences $(\overline{a}_\alpha)_{\alpha<\kappa}$ where $\overline{a}_\alpha = \langle a_{\alpha,i} : i < \lambda\rangle$ is called a *mutually indiscernible array* over a set of parameters C if, for each $\alpha < \kappa$, the sequence $\overline{a}_\alpha$ is an indiscernible sequence over $C\overline{a}_{\neq\alpha}$.

The following two lemmas are essentially [Adl14, Lemma 8]. We include a proof for completeness.

Lemma 3.3.9. *If $\overline{a} = (a_i)_{i<\omega}$ is an M-indiscernible sequence, there is an indiscernible global M-invariant type $q \supseteq tp(\overline{a}/M)$.*

Proof. Let N be an $|M|^+$-saturated elementary extension of M of size κ and let $r \supseteq \text{tp}(\overline{a}/M)$ be an arbitrary global M-invariant type extending $\text{tp}(\overline{a}/M)$. Let $\overline{b} = (b_i)_{i<\omega} \models r|_N$. By Ramsey and compactness, we may extract from $\overline{b}$ an N-indiscernible sequence $(c_i)_{i<\omega}$. Clearly $\text{tp}(\overline{c}/N)$ extends $\text{tp}(\overline{a}/M)$. It is also M-invariant: if not, there are $n \equiv_M n'$ in N, an increasing k-tuple $\overline{i}$ from ω, and a formula φ so that

$$\models \varphi(c_{\overline{i}}; n) \leftrightarrow \neg\varphi(c_{\overline{i}}, n').$$

Then there is an increasing k-tuple $\overline{j}$ so that

$$\models \varphi(b_{\overline{j}}; n) \leftrightarrow \neg\varphi(b_{\overline{j}}; n'),$$

since the sequence $\overline{c}$ is extracted from $\overline{b}$. This contradicts the fact that $\overline{b}$ realizes an M-invariant type over N. By Fact 3.3.2, the type $\text{tp}(\overline{c}/N)$ determines a unique M-invariant extension to $\mathbb{M}$. Call it q. Then q is an indiscernible type. $\square$

Lemma 3.3.10. *Suppose $M \models T$, $\overline{a} = (a_i)_{i<\omega}$ is an M-indiscernible sequence, and $q \supseteq$ tp$(\overline{a}/M)$ is a global M-invariant indiscernible type. Let $(\overline{a}_i)_{i<\omega} \models q^{\otimes\omega}|_M$ with $\overline{a}_0 = \overline{a}$, where $\overline{a}_i = (a_{i,j})_{j<\omega}$. Then $(\overline{a}_i)_{i<\omega}$ is a mutually indiscernible array over M.*

Proof. We prove by induction on n that $(\overline{a}_i)_{i\leq n}$ is mututally indiscernible over M. For $n = 1$, there's nothing to prove. Suppose it's been shown for n and consider $(a_i)_{i\leq n+1}$. As q is an indiscernible type, $\overline{a}_{n+1}$ is $M\overline{a}_{\leq n}$-indiscernible. For $i \leq n$, we know, by induction, that $\overline{a}_i$ is $M\overline{a}_{<i}\overline{a}_{i+1}\ldots\overline{a}_n$-indiscernible. As $\overline{a}_{n+1} \models q|_{M\overline{a}_{\leq n}}$, this entails $\overline{a}_i$ is indiscernible over $M\overline{a}_{<i}\overline{a}_{i+1}\ldots\overline{a}_{n+1}$, which completes the induction. $\square$

Kim-dividing

In this subsection, we define Kim-dividing and Kim-forking, the fundamental notions explored in this paper. To start, we will need the definition of q-dividing, introduced by Hrushovski in [Hru12, Section 2.1]:

Definition 3.3.11. Suppose $q(y)$ is an A-invariant global type. The formula $\varphi(x; y)$ q-divides over A if for some (equivalently, any) Morley sequence $\langle b_i : i < \omega\rangle$ in q over A, $\{\varphi(x; b_i) : i < \omega\}$ is inconsistent.

We note that we will consistently use the letters p, q, r to refer to types, n, m, k, l to refer to numbers. In this way, no confusion between q-dividing and the more familiar k-dividing will arise.

The related notion of a *higher formula* was introduced by Malliaris and Shelah in [MS15a] on the way to a new characterization of NTP_1 theories:

Definition 3.3.12. [MS15a, Definition 8.6] A *higher formula* is a triple $(\varphi, A, \mathcal{D})$ where $\varphi = \varphi(x; y)$ is a formula, A is a set of parameters, and $\mathcal{D}$ is an ultrafilter on $A^{l(y)}$ so that, if $q = \text{Av}(\mathcal{D}, \mathbb{M})$ and $\langle b_i : i < \omega\rangle \models q^{\otimes\omega}|_A$ then $\{\varphi(x; b_i) : i < \omega\}$ is consistent.

We can rephrase the above definition as: $(\varphi, A, \mathcal{D})$ is a higher formula if, setting $q = \mathrm{Av}(\mathcal{D}, \mathbb{M})$, $\varphi(x; y)$ does not q-divide over A.

Definition 3.3.13. We say that a formula $\varphi(x; b)$ *Kim-divides* over A if there is some A-invariant global type $q \supseteq \mathrm{tp}(b/A)$ so that $\varphi(x; y)$ q-divides. The formula $\varphi(x; b)$ *Kim-forks* over A if $\varphi(x; b) \vdash \bigvee_{i<k} \psi_i(x; c^i)$ and each $\psi_i(x; c^i)$ Kim-divides over A. A type Kim-forks if it implies a formula which does. If $\mathrm{tp}(a/Ab)$ does not Kim-fork over A, we write $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_A b$.

We call this notion *Kim*-dividing to make explicit the fact that this definition was inspired by a suggestion of Kim in his 2009 BIRS talk [Kim09], where he proposed an independence relation based on instances of dividing that are witnessed by every appropriate Morley sequence. A rough connection between Kim's notion and ours is provided by Theorem 3.3.16 below, which shows that, in an NSOP_1 theory, dividing with respect some invariant Morley sequence is equivalent to dividing with respect to all. An even tighter connection is established by Theorem 3.7.7, which shows that we can drop the assumption that the Morley sequences are generated by an invariant type. (We note that for technical reasons our notion is still different from Kim's – the proposal of [Kim09] forces a kind of base monotonicity and we do not).

In general, we only know that a type over A has a global A-invariant extension when A is a model. Thus, when working with Kim-independence below, we will restrict ourselves almost entirely to the case where the base is a model.

The next two propositions explain how the notions of higher formula and q-dividing interact with SOP_1.

Proposition 3.3.14. *Suppose T has SOP_1. Then there is a model $M \models T$, a formula $\varphi(x; b)$, and ultrafilters $\mathcal{D}_0, \mathcal{D}_1$ on M with*

$$Av(\mathcal{D}_0, M) = Av(\mathcal{D}_1, M) = tp(b/M),$$

so that $(\varphi, M, \mathcal{D}_0)$ is higher but $(\varphi, M, \mathcal{D}_1)$ is not higher.

Proof. Fix a Skolemization T^{Sk} of M. As T has SOP_1, there is, by Proposition 5.2.1, a formula $\varphi(x; y)$ and an array $(c_{i,j})_{i<\omega+1, j<2}$ such that

1. $(\bar{c}_i)_{i<\omega+1}$ is an indiscernible sequence (with respect to the Skolemized language)
2. $c_{\omega,0} \equiv^{L^{Sk}}_{\bar{c}_{<\omega}} c_{\omega,1}$.
3. $\{\varphi(x; c_{i,0}) : i < \omega + 1\}$ is consistent.
4. If $i < j$, then $\{\varphi(x; c_{i,1}), \varphi(x; c_{j,1})\}$ is inconsistent.

Put $M = \mathrm{Sk}(\bar{c}_{<\omega})$. For $j = 0, 1$, let $\mathcal{D}_j$ be any non-principal ultrafilter on M, concentrating on $\langle c_{i,j} : i < \omega \rangle$ and set $q_j = \mathrm{Av}_L(\mathcal{D}_j, \mathbb{M})$ for $j = 0, 1$. Note that $q_0|_M = \mathrm{tp}_L(c_{\omega,0}/M) = \mathrm{tp}_L(c_{\omega,1}/M) = q_1|_M$ by (2). By (3), $\varphi(x; y)$ does not q_0-divide, hence $(\varphi, M, \mathcal{D}_0)$ is higher. However, by (4), $\{\varphi(x; c_{j,1}) : j < \omega\}$ is 2-inconsistent hence $\varphi(x; y)$ q_1-divides, so $(\varphi, M, \mathcal{D}_1)$ is not higher. $\square$

Proposition 3.3.15. *Suppose A is a set of parameters and $\varphi(x;b)$ is a formula which q-divides over A for some global A-invariant type $q \supseteq tp(b/A)$. If there is some global A-invariant $r \supseteq tp(b/A)$ such that $\varphi(x;y)$ does not r-divide, then T has SOP_1.*

Proof. As $\varphi(x;y)$ q-divides over A, there is k so that instances of $\varphi(x;y)$ instantiated on a Morley sequence of q are k-inconsistent.

Let $(c_{i,1}, c_{i,0})_{i\in\mathbb{Z}} \models (q\otimes r)^{\otimes\mathbb{Z}}|_M$. We have to check that the sequence satisfies the following properties:

1. $\{\varphi(x;c_{i,0}) : i \in \mathbb{Z}\}$ is consistent
2. $\{\varphi(x;c_{i,1}) : i \in \mathbb{Z}\}$ is k-inconsistent
3. $c_{i,0} \equiv_{\overline{c}_{>i}} c_{i,1}$ for all $i \in \mathbb{Z}$.

Note that $(c_{i,0})_{i\in\mathbb{Z}} \models r^{\otimes\mathbb{Z}}|_M$ so (1) follows from our assumption that $\varphi(x;y)$ does not r-divide. Likewise, $(c_{i,1})_{i\in\mathbb{Z}} \models q^{\otimes\mathbb{Z}}|_M$ so (2) follows from the fact that $\varphi(x,y)$ q-divides. Finally, for any $i \in \mathbb{Z}$, we have $\overline{c}_{>i}$ realizes a global M-invariant type over $Mc_{i,0}c_{i,1}$. Hence (3) follows from the fact that $c_{i,0} \equiv_M c_{i,1}$. □

Theorem 3.3.16. *The following are equivalent for the complete theory T:*

1. *T is $NSOP_1$*
2. *Ultrafilter independence of higher formulas: for every model $M \models T$, and ultrafilters $\mathcal{D}$ and $\mathcal{E}$ on M with $Av(\mathcal{D}, M) = Av(\mathcal{E}, M)$, $(\varphi, M, \mathcal{D})$ is higher if and only if $(\varphi, M, \mathcal{E})$ is higher.*
3. *Kim's lemma for Kim-dividing: For every model $M \models T$ and $\varphi(x;b)$, if $\varphi(x;y)$ q-divides for some global M-invariant $q \supseteq tp(b/M)$, then $\varphi(x;y)$ q-divides for* every *global M-invariant $q \supseteq tp(b/M)$.*

Proof. (1) $\Longrightarrow$ (3) is the contrapositive of Proposition 3.3.15.

(2) $\Longrightarrow$ (1) is the contrapositive of Proposition 3.3.14.

(3) $\Longrightarrow$ (2): Immediate, since every type finitely satisfiable in M is M-invariant. □

Remark 3.3.17. Note that the proof gives a bit more: if T is NSOP_1, (2) is true over arbitrary sets and (3) is true over an arbitrary set A as well, though this may be vacuous if $\mathrm{tp}(b/A)$ does not extend to a global A-invariant type.

The basic properties of Kim-independence

Theorem 3.3.16, a kind of Kim's lemma for Kim-dividing, already gives a powerful tool for proving that in NSOP$_1$ theories Kim-independence enjoys many of the properties known to hold for non-forking independence in simple theories.

We will frequently use the following easy observation. The proof is exactly as in the case of dividing. See, e.g., [GIL02, Lemma 1.5] or [She80, Lemma 1.4].

Lemma 3.3.18. *(Basic Characterization of Kim-dividing) Suppose T is an arbitrary complete theory. The following are equivalent:*

1. *$tp(a/Ab)$ does not Kim-divide over A.*
2. *For any global A-invariant $q \supseteq tp(b/A)$ and $I = \langle b_i : i < \omega \rangle \models q^{\otimes \omega}|_A$ with $b_0 = b$, there is $a' \equiv_{Ab} a$ such that I is Aa'-indiscernible.*
3. *For any global A-invariant $q \supseteq tp(b/A)$ and $I = \langle b_i : i < \omega \rangle \models q^{\otimes \omega}|_A$ with $b_0 = b$, there is $I' \equiv_{Ab} I$ such that I' is Aa-indiscernible.*

Note that in an NSOP$_1$ theory, by Kim's Lemma for Kim-dividing, we could have replaced (2) by: *there is* a global A-invariant $q \supseteq \text{tp}(b/A)$ and $I = \langle b_i : i < \omega \rangle \models q^{\otimes \omega}|_A$ with $b_0 = b$, so that for some $a' \equiv_{Ab} a$ I is Aa'-indiscernible (and similarly for (3)), provided $\text{tp}(b/A)$ extends to a global A-invariant type.

The following proposition is proved by the same argument one uses to prove forking = dividing via Kim's lemma, as in [GIL02, Theorem 2.5] or [CK12, Corollary 3.16].

Proposition 3.3.19. *(Kim-forking = Kim-dividing) Suppose T is NSOP$_1$. If $M \models T$, if $\varphi(x; b)$ Kim-forks over M then $\varphi(x; b)$ Kim-divides over M.*

Proof. Suppose $\varphi(x; b) \vdash \bigvee_{j<k} \psi_j(x; c^j)$ where each $\psi_i(x; c^i)$ Kim-divides over M. Fix an ultrafilter $\mathcal{D}$ on M so that $(b, c^0, \ldots, c^{k-1}) \models \text{Av}(\mathcal{D}, M)$. Let $(b_i, c_i^0, \ldots, c_i^{k-1})_{i<\omega}$ be a Morley sequence in $\text{Av}(\mathcal{D}, \mathbb{M})$. Then $(b_i)_{i<\omega}$ is an M-invariant Morley sequence. We must show $\{\varphi(x; b_i) : i < \omega\}$ is inconsistent. Suppose not – let $a \models \{\varphi(x; b_i) : i < \omega\}$. We have $\varphi(x; b_i) \vdash \bigvee_{j<k} \psi_j(x; c_i^j)$ so for each $i < \omega$, there is $j(i) < k$ so that $\models \psi_{j(i)}(a; c_i^{j(i)})$. By the pigeonhole principle, there is $j_* < k$ so that $X = \{i < \omega : j(i) = j_*\}$ is infinite. Then $(c_i^{j_*})_{i \in X}$ is an M-invariant Morley sequence in $\text{tp}(c^{j_*}/M)$. As T is NSOP$_1$, Kim-dividing over M is witnessed by any M-invariant Morley sequence so $\{\psi_{j_*}(x; c_i^{j_*}) : i \in X\}$ is inconsistent. But $a \models \{\psi_{j_*}(x; c_i^{j_*}) : i \in X\}$, a contradiction. $\square$

Proposition 3.3.20. *(Extension over Models) Suppose M is a model, and $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$. Then for any c, there is $a' \equiv_{Mb} a$ so that $a' \mathop{\smile\!\!\!\!\!\!|}^K_M bc$.*

Proof. This is exactly as in the usual proof that forking satisfies extension. Let $p(x; b) = \text{tp}(a/Mb)$. We claim that the following set of formulas is consistent:

$$p(x; b) \cup \{\neg\psi(x; b, c) : \psi(x; b, c) \in L(Mbc) \text{ and } \psi(x; b, c) \text{ Kim-divides over } M\}.$$

If this set of formulas is not consistent, then by compactness,

$$p(x;b) \vdash \bigvee_{i<k} \psi_i(x;b,c_i),$$

where each $\psi_i(x;b,c_i)$ Kim-divides over M. It follows that $\text{tp}(a/Mb)$ Kim-forks over M, a contradiction. So this set is consistent and we may choose a realization a'. Then $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M bc$ and $a' \equiv_{Mb} a$. □

Proposition 3.3.21. *(Chain Condition for Invariant Morley Sequences) Suppose T is NSOP$_1$ and $M \models T$. If $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$ and $q \supseteq tp(b/M)$ is a global M-invariant type, then for any $I = \langle b_i : i < \omega \rangle \models q^{\otimes\omega}|_M$ with $b = b_0$, there is $a' \equiv_{Mb} a$ so that $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits_M I$ and I is Ma'-indiscernible.*

Proof. By the basic characterization of Kim-dividing, Lemma 3.3.18, given $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$, $q \supseteq \text{tp}(b/M)$ a global M-invariant type, and $I = \langle b_i : i < \omega \rangle \models q^{\otimes\omega}|_M$ with $b = b_0$, there is $a' \equiv_{Mb} a$ so that I is Ma'-indiscernible. To prove the proposition it suffices to show $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b_{<n}$ for all n. Given $n < \omega$, let $r(x;y_0,\ldots,y_{n-1}) = \text{tp}(a';b_0,\ldots,b_{n-1}/M)$. Then $\langle (b_{kn+n-1}, b_{kn+n-2},\ldots,b_{kn}) : k < \omega \rangle \models (q^{\otimes n})^{\otimes\omega}|_M$ and, by indiscernibility,

$$a' \models \bigcup_{k<\omega} r(x;b_{kn+n-1}, b_{kn+n-2},\ldots,b_{kn}).$$

As T is NSOP$_1$, this shows $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b_{<n}$. □

Section 3.5 will be dedicated to the proof that $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K$ is symmetric in NSOP$_1$ theories. The argument will require more tools, but at this stage we can already observe the converse: even a weak form of symmetry for $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K$ will imply that a theory is NSOP$_1$.

Proposition 3.3.22. *The following are equivalent for a complete theory T:*

1. *T is NSOP$_1$.*
2. *Weak symmetry: if $M \models T$, then $b \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^i_M a \implies a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$.*

Proof. (1) $\implies$ (2). Suppose T is NSOP$_1$. As $b \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^i_M a$, there is a global M-invariant type $r \supseteq \text{tp}(b/Ma)$. We can find a Morley sequence $I = \langle b_i : i < \omega \rangle$ in $q|_{Ma}$ with $b_0 = b$. Then I is Ma-indiscernible, so no formula in $\text{tp}(a/Mb)$ divides with respect to the sequence I. But by Kim's lemma for Kim-dividing, this implies $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$.

(2) $\implies$ (1). Suppose T has SOP$_1$. Then by Theorem I.5.1, there is a model $M \models T$, $a_0b_0 \equiv_M a_1b_1$ with $b_i \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^i_M a_i$ and $b_1 \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^i_M b_0$, but, setting $p(x;b_0) = \text{tp}(a_0/Mb_0)$, we have $p(x;b_0) \cup p(x;b_1)$ is inconsistent. As $b_0 \equiv_M b_1$ and $b_1 \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^i_M b_0$, (b_0,b_1) starts a Morley sequence in some M-invariant type, $\langle b_i : i < \omega \rangle$. As $\bigcup_{i<\omega} p(x;b_i)$ is inconsistent, we have $a_0 \not\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b_0$. Since $b_0 \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^i_M a_0$, weak symmetry fails. □

3.4 Local character

In this section, we prove local character for $\mathop{\smile\hskip-0.9em\raise0.7ex\hbox{$\mid$}}^K$ in NSOP_1 theories: for every NSOP_1 theory, there is a cardinal κ so that, given a model $M \models T$ and a type $p \in S(M)$, there is an elementary submodel $M' \preceq M$ of size $< \kappa$ such that p does not Kim-fork over M'. We give a simple and soft argument showing first that κ can be taken to be the first measurable cardinal above $|T|$. Then, in a more difficult argument, we show that κ can be taken to be $(2^{|T|})^+$. The argument involving large cardinals is, of course, implied by the stronger result, but we thought that the conceptual simplicity of the first argument might be helpful in understanding the second. Lastly, we show that for any regular κ, we can construct a model which satisfies local character—this clarifies the situation for cardinals between $|T|$ and $2^{|T|}$.

In order to prove our first theorem, we will use the following facts about measurable cardinals:

Fact 3.4.1. [Kan03, Theorem 7.17] Suppose that $\mu > |T|$ is a measurable cardinal and that $\mathcal{U}$ is a normal (non-principal) ultrafilter on μ. Suppose that $(a_i)_{i<\mu}$ is a sequence of finite tuples in $\mathbb{M}$, then for some set $X \in \mathcal{U}$, $(a_i)_{i \in X}$ is an indiscernible sequence.

Fact 3.4.2. [KLS16, Fact 2.9] If $A = \bigcup_{i<\mu} A_i \subseteq \mathbb{M}$ is a continuous increasing union of sets where $|A_i| < \mu$, $B \subseteq \mathbb{M}$ is some set of cardinality $< \mu$, and $(a_i)_{i<\mu}$, $\mathcal{U}$ are as in Fact 3.4.1 with a_i tuples from A, then for some set $X \in \mathcal{U}$, $(a_i)_{i \in X}$ is fully indiscernible over B (with respect to A and $(A_i)_{i<\mu}$), which means that for every $i \in X$ and $j < i$ in X, we have $a_j \subseteq A_i$, and $(a_j)_{i \leq j \in X}$ is indiscernible over $A_i \cup B$.

Theorem 3.4.3. *Suppose that T is $NSOP_1$ and that $|T| < \mu$ is measurable. Suppose that $M \models T$. Then for every $p \in S(M)$ there is a model $N \prec M$ with $|N| < \mu$ such that p does not Kim-fork over N.*

Proof. Suppose not. Construct by induction on $i < \mu$ a sequence $((M_i, b_i, b'_i, \varphi_i))_{i<\mu}$ such that:

- $(M_i)_{i<\mu}$ is an increasing continuous sequence of models.
- For $i < \mu$, $\varphi_i(x, y_i, y'_i)$ is a formula in L.
- $b'_i \in M_i$ and $b_i \in M \setminus M_i$.
- $\varphi(x, b_i, b'_i) \in p$ witnesses Kim-dividing of p over M_i.
- For $i < \mu$, M_i is some model containing $\{b_j : j < i\}$ of size $|T| + |i|$.

We can construct such a sequence by our assumption.

Note that all clubs $E \subseteq \mu$ are in $\mathcal{U}$ (see the proof of Fact 3.4.2 in [KLS16, Fact 2.9]).

By Fodor's lemma for normal ultrafilters [Kan03, Exercise 5.10], applied to the function $F : \lim(\mu) \to \mu$ such that $F(\delta) = \min\{i < \delta : b'_\delta \in M_i\}$, there is some $X_0 \in \mathcal{U}$ (consisting of limit ordinals) such that for all $i \in X_0$, $b'_i = b'$ is constant. For convenience assume that $b' \in M_0$.

By Fact 3.4.2, there is some $X_1 \subseteq X_0$ in $\mathcal{U}$ such that $(b_i)_{i\in X_1}$ is fully indiscernible with respect to $(M_i)_{i<\mu}$.

Let $X_2 \subseteq X_1$ in $\mathcal{U}$ be such that for all $i \in X_2$ if $j < i$ then there is $j < \alpha < i$ such that $\alpha \in X_1$ (as X_1 is unbounded, the set of all $i < \mu$ such that for all $j < i$ there is such an α is a club, so in $\mathcal{U}$).

Let $\alpha_0 \in X_2$. Then $(b_i)_{i\geq\alpha_0, i\in X_2}$ is an indiscernible sequence such that $b_i \mathop{\smile\!\!\!\!\!\!\perp}^u_{M_{\alpha_0}} b_{>i}$ for all $\alpha_0 \leq i \in X_2$. To see this, suppose that $\psi(b_i, b_{>i}, m)$ holds where $m \in M_{\alpha_0}$. Then $m \in M_\beta$ for some $\beta < \alpha_0$ (as α_0 is limit). So by definition of X_2, there is some $\beta < j < \alpha_0$ in X_1. But then $(b_\varepsilon)_{\varepsilon\geq j}$ is indiscernible over M_β by choice of X_1, so $\psi(b_j, b_{>i}, m)$ holds as well. But $b_j \in M_{\alpha_0}$ by construction.

We get a contradiction, since $\{\varphi(x, b_i, b') : i \in X_2, i \geq \alpha_0\}$ is in p but also inconsistent since $\varphi(x, b_{\alpha_0}, b')$ Kim-divides over M_{α_0}. □

Lemma 3.4.4. *Suppose that N is some model and that $p \in S(\mathbb{M})$ is a global type finitely satisfiable in N which extends $tp(c/N)$. Given any set $A \subseteq N$, there is some $B \prec N$ of size $\leq |T| + |A|$ such that $A \subseteq B$ and $p^{\otimes\omega}|_B$ is a type of a Morley sequence generated by some global type finitely satisfiable in B.*

In particular, if $\varphi(x, c)$ Kim-divides over N then $\varphi(x, c)$ Kim-divides over B.

Proof. Let $p \in S(\mathbb{M})$ be a global type extending $\mathrm{tp}(c/N)$, finitely satisfiable in N.

Let B_0 be any model containing A of size $\leq |A|+|T|$, and let $\bar{c} \models p^{\otimes\omega}|_N$. Let $N \supseteq B'_1 \supseteq B_0$ be such that for every $n < \omega$ and every formula $\psi(y, c_{<n})$ over B_0, if $\mathbb{M} \models \psi(c_n, c_{<n})$ then $\psi(y, c_{<n})$ is satisfiable in B'_1 and let B_1 be any model containing B'_1 of size $|T|+|A|$. Continue like this, and finally, let $B = \bigcup_{i<\omega} B_i$. Then $\bar{c}$ is still a Morley sequence sequence over B in a B-finitely satisfiable type (note that it is indiscernible). □

Theorem 3.4.5. *Suppose T is NSOP$_1$. Then for any $M \models T$ and $p \in S(M)$, there is $M' \preceq M$ so that p does not Kim-fork over M' and $|M'| \leq 2^{|T|}$.*

Proof. Let $\kappa = (2^{|T|})^+$—κ is a regular cardinal, greater than $2^{|T|}$, and $\mu < \kappa$ implies $\mu^{|T|} < \kappa$ (these are the only properties of κ we will use).

Suppose not. Then there is some $p \in S(M)$ witnessing this. Clearly $|M| \geq \kappa$. For every $i < \kappa$ we can find c_i, d_i, N_i, and $\varphi_i(x, y_i, z_i)$ such that:

- $c_i \in N_{i+1} \setminus N_i$, $d_i \in N_i$, $\langle N_i : i < \kappa\rangle$ is increasing continuous, $|N_i| \leq |T| + |i| < \kappa$, $\varphi_i(x, y_i, z_i)$ is a formula such that $\varphi_i(x, c_i, d_i)$ Kim-divides over N_i and is in p.

Let S be $\{\delta < \kappa : \mathrm{cof}(\delta) = |T|^+\}$. Then S is a stationary set.

For every $\delta \in S$, fix some global coheir $q_\delta \in S(\mathbb{M})$ over N_δ extending $\mathrm{tp}(c_\delta/N_\delta)$. Given a partition of a stationary subset of κ into $< \kappa$ parts, one of these has to be a stationary set. Hence, we may assume that for every $\delta \in S$, $\varphi_\delta = \varphi$ and $\varphi(x, c_i, d_i)$ is k-Kim-dividing for some fixed k, witnessed by any Morley sequence in q_δ. Define the regressive function

$f : S \to \kappa$ by $f(\delta) = \min\{i < \delta : d_\delta \in N_i\}$ (this set is non-empty by continuity of the sequence). By Fodor's lemma, we may assume that f is constant on S, and further restricting it, we may even assume that $d_\delta = d$ is fixed for every $\delta \in S$. This allows us to assume for simplicity that $d = \emptyset$.

By Lemma 3.4.4, for every $\delta \in S$ there is some $M_\delta \prec N_\delta$ of size $|T|$ such that $\varphi(x, c_\delta)$ Kim-divides over M_δ, and moreover, such that $q_\delta^{\otimes\omega}|_{M_\delta}$ is a type of a Morley sequence of some global coheir r_δ over M_δ.

As $\operatorname{cof}(\delta) = |T|^+$ for every $\delta \in S$, for each such δ there is some $i < \delta$ such that $M_\delta \prec N_i$. Hence by Fodor's lemma, there is some $i < \kappa$ and a stationary $S' \subseteq S$ such that for every $\delta \in S'$, $M_\delta \prec N_i$. Then we can find some model M_0^*, a global coheir r_0^* over M_0^* and a stationary $S_0 \subseteq S'$ such that for every $\delta \in S_0$, $M_\delta = M_0^*$ (note that the number of possible M_δ's is $\leq |N_i|^{|T|} < \kappa$) and $q_\delta^{\otimes\omega}|_{M_0^*} = r_\delta^{\otimes\omega}|_{M_0^*} = r_0^{*\otimes\omega}|_{M_0^*}$ (the number of ω-types over M_0^* is $\leq 2^{|M_0^*|} < \kappa$ as $|M_0^*| = |T|$).

Let $\delta_0 = \min S_0$ and $e_0 = c_{\delta_0}$.

By Lemma 3.4.4, for every $\delta \in S_0 \setminus \{\delta_0\}$ there is some $M_0^* c_{\delta_0} \subseteq M_\delta \prec N_\delta$ of size $|T|$ such that $\varphi(x, c_\delta)$ Kim-divides over M_δ, and moreover, such that $q_\delta^{\otimes\omega}|_{M_\delta}$ is a type of a Morley sequence of some global coheir over M_δ. Thus, as above, we can find some stationary $S_1 \subseteq S_0 \setminus \{\delta_0\}$, M_1^* and r_1^* such that for every $\delta \in S_1$, $M_\delta = M_1^*$ and $q_\delta^{\otimes\omega}|_{M_1^*} = r_1^{*\otimes\omega}|_{M_1^*}$. Let $\delta_1 = \min S_1$ and $e_1 = c_{\delta_1}$.

Continuing like this we find and increasing sequence $\langle \delta_i : i < \omega \rangle$ of ordinals in κ, an increasing sequence of models $\langle M_i^* : i < \omega \rangle$, $e_i \in M$ for $i < \omega$ and global coheirs (over M_i^*) r_i^* such that:

- M_i^* contains $e_{<i}$, $\varphi(x, e_j)$ is k-Kim-dividing over M_i^* for every $i < j$, r_i^* is a global coheir over M_i^* such that for all $i \leq j$, r_i^* extends $\operatorname{tp}(e_j / M_i^*)$ (in particular, $e_j \equiv_{M_i^*} e_i$ for all $j \geq i$) and $r_j^{*\otimes\omega}|_{M_i^*} = r_i^{*\otimes\omega}|_{M_i^*}$.

Denote $\bar{e} = \langle e_i : i < \omega \rangle$. Note that $\{\varphi(x; e_i) : i < \omega\}$ is a subset of p, hence consistent.

Claim: Suppose $i_0 < \ldots < i_{n-1} < \omega$ and for each $j < n$, $f_j \models r_{i_j}^*|_{M_{i_j}^* \bar{e} f_{>j}}$. Then

1. $e_{i_j} \equiv_{e_{i_{<j}} f_{<j}} f_j$ for all $j < n$
2. $\{\varphi(x; f_j) : j < n\}$ is k-inconsistent.

Proof of claim: By induction on n, we prove that if $i_0 < \ldots < i_{n-1} < \omega$, then $e_{i_j} \equiv_{M_{i_0}^* e_{<i_j} f_{<j}} f_j$. For $n = 0$ there is nothing to prove. Suppose the claim is true for n and we are given $i_0 < \ldots < i_n$ and $(f_j)_{j<n+1}$ with $f_j \models r_{i_j}^*|_{M_{i_j}^* \bar{e} f_{>j}}$ for all $j < n+1$. Then clearly $f_0 \equiv_{M_{i_0}^*} e_{i_0}$. For $0 < j < n+1$, by induction $f_j \equiv_{M_{i_1}^* e_{<i_j} f_{\in[1,j)}} e_{i_j}$, hence $f_j \equiv_{M_{i_0}^* e_{<i_j} f_{\in[1,j)}} e_{i_j}$. As $f_0 \mathop{\smile\!\!\!\!\!\mid}^u_{M_{i_0}^*} \bar{e} f_{>0}$, we get that $f_j \equiv_{M_{i_0}^* e_{<i_j} f_{<j}} e_{i_j}$. This shows (1). To see (2), note that $(f_{n-1}, f_{n-2}, \ldots, f_0) \models r_{i_0}^{\otimes n}|M_{i_0}^* = q_{\delta_{i_0}}^{\otimes n}|M_0^*$, hence $\{\varphi(x; f_j) : j < n\}$ is k-inconsistent, by our assumption that $\varphi(x; c_{\delta_0})$ k-Kim-divides with respect to Morley sequences in q_{δ_0}. $\square$

By compactness, we can find an array $(c_{i,0}, c_{i,1})_{i<\omega}$ so that $\{\varphi(x; c_{i,0}) : i < \omega\}$ is consistent, $\{\varphi(x; c_{i,1}) : i < \omega\}$ is k-inconsistent, and $c_{i,0} \equiv_{\bar{c}_{<i}} c_{i,1}$ for all $i < \omega$. By Lemma 4.2.2, we obtain SOP_1, a contradiction. $\square$

Corollary 3.4.6. *Suppose that T is a complete theory. The following are equivalent.*

1. *For some uncountable cardinal κ, there is no sequence $\langle N_i, \varphi_i(x, y_i), c_i : i < \kappa \rangle$ such that $\langle N_i : i < \kappa \rangle$ is an increasing continuous sequence of models of T, $\varphi_i(x, y_i)$ is a formula over N_i, $c_i \in N_{i+1}$, such that $\varphi_i(x, c_i)$ Kim-forks over N_i and $\{\varphi(x, c_i) : i < \kappa\}$ is consistent.*

2. *T is $NSOP_1$.*

Proof. (2) implies (1) by the proof of Theorem 3.4.5 (with $\kappa = (2^{|T|})^+$).

(1) implies (2). This is a variation on the proof of Proposition 3.3.14. Suppose T has SOP_1 as witnessed by some formula $\varphi(x, y)$. Let T^{sk} be a Skolemized expansion of T. Then T^{sk} also has SOP_1 as witnessed by $\varphi(x, y)$. Thus by Proposition 2.4, we can find a formula $\varphi(x, y)$ and an array $(c_{i,j})_{i<\omega, j<2}$ such that $c_{i,0} \equiv_{\bar{c}_{<i}} c_{i,1}$ for all $i < \omega$, $\{\varphi(x, c_{i,0}) : i < \omega\}$ is consistent and $\{\varphi(x, c_{i,1}) : i < \omega\}$ is 2-inconsistent (all in $\mathbb{M}^{sk}$). By Ramsey and compactness we may assume that $\langle \bar{c}_i : i < \omega \rangle$ is indiscernible (with respect to $\mathbb{M}^{sk}$). Extend this sequence to one of length κ.

For $i < \kappa$, let $N_i = \mathrm{dcl}(\bar{c}_{<i})$ (in $\mathbb{M}^{sk}$). Then for every limit ordinal $\delta < \kappa$, $\varphi(x, c_{\delta,1})$ Kim-divides over N_δ as the sequence $\langle c_{j,1} : \delta \leq j < \kappa \rangle$ is indiscernible and for all $\delta \leq j$, $\bar{c}_j \mathop{\smile\!\!\!\!\!\!\mid}^u_{N_\delta} \bar{c}_{>j}$. As $c_{\delta,1} \equiv_{\bar{c}_{<\delta}} c_{\delta,0}$, it follows that $c_{\delta,1} \equiv_{N_\delta} c_{\delta,0}$, and hence $\varphi(x, c_{\delta,0})$ also Kim-divides. As κ is uncountable, $\mathrm{otp}(\lim(\kappa)) = \kappa$, so $\langle N_\delta, \varphi(x, y), c_{\delta,0} : \delta < \kappa \rangle$ contradicts (1). $\square$

Theorem 3.4.7. *Suppose that T is $NSOP_1$. Then for every regular cardinal $\kappa > |T|$ there is a model M of size κ such that for all $p \in S(M)$ there is $N \prec M$ with $|N| < \kappa$ such that p does not Kim-fork over N.*

Proof. Let $I = (a_i)_{i<\kappa}$ be an indiscernible sequence with respect to T^{sk} — a Skolemized expansion of T. Let $M = \mathrm{dcl}(I)$. Let $p \in S(M)$. For $i < \kappa$ let

$$N_i = \mathrm{dcl}((a_j)_{j<i}).$$

Suppose for contradiction that for every $i < \kappa$, p Kim-forks over N_i.

This means that for every $i < \kappa$ there is a formula $\varphi_i(x, t_i(b_i, b_i'))$ witnessing Kim-dividing over N_i, where t_i is a Skolem term, $b_i \subseteq \{a_j : j \geq i\}$, $b_i' \subseteq \{a_j : j < i\}$, and both are increasing tuples.

Let $E \subseteq \kappa$ be the set of limits $\alpha \in \kappa$ such that for all $i < \alpha$, $b_i \subseteq (a_j)_{j<\alpha}$. Then E is a club of κ. Define $F : E \to \kappa$ by $F(\alpha) = \max\{j : a_j \in b_\alpha'\}$. By Fodor's lemma there is a stationary set $S \subseteq E$ on which F is constant γ. Reducing to an unbounded subset of S, we may assume that for every $\alpha \in S$, $\varphi_\alpha = \varphi$, $t_\alpha = t$ and $b_\alpha' = b'$ (all the b_α' come from $\{a_j : j < \gamma\}$ which has size $|\gamma| < |S| = \kappa$). By choice of E, for all $\alpha < \beta$ from S, $b_\alpha < b_\beta$ (i.e., every coordinate of b_β is greater than every coordinate of b_α). Hence $(t(b_\alpha, b'))_{\alpha_0 \leq \alpha \in S}$ is an indiscernible sequence over N_{α_0} such that $t(b_\alpha, b') \mathop{\smile\!\!\!\!\!\!\mid}^u_{N_{\alpha_0}} t(b_{>\alpha}, b')$ for every $\alpha_0 \in S$ by the construction of N_{α_0}, and hence, as $\varphi(x, t(b_{\alpha_0}, b'))$ Kim-divides over N_{α_0}, $\{\varphi(x, t(b_\alpha, b')) : \alpha_0 \leq \alpha \in S\}$ in inconsistent, but it is contained in p. □

Remark 3.4.8. Note that this theorem is most interesting for the case $|T| < \kappa \leq 2^{|T|}$, as this is not covered by Theorem 3.4.5.

3.5 Symmetry

Generalized indiscernibles and a class of trees

For an ordinal α, let the language $L_{s,\alpha}$ be $\langle \trianglelefteq, \wedge, <_{lex}, (P_\beta)_{\beta<\alpha} \rangle$. We may view a tree with α levels as an $L_{s,\alpha}$-structure by interpreting $\trianglelefteq$ as the tree partial order, $\wedge$ as the binary meet function, $<_{lex}$ as the lexicographic order, and P_β interpreted to define level β. For the rest of the paper, a tree will be understood to be an $L_{s,\alpha}$-structure for some appropriate α. We will sometimes suppress the α and refer instead to L_s, where the number of predicates is understood from context. We define a class of trees $\mathcal{T}_\alpha$ as follows.

Definition 3.5.1. Suppose α is an ordinal. We define $\mathcal{T}_\alpha$ to be the set of functions f so that

- $\mathrm{dom}(f)$ is an end-segment of α, possibly empty unless α is a limit.
- $\mathrm{ran}(f) \subseteq \omega$.
- finite support: the set $\{\gamma \in \mathrm{dom}(f) : f(\gamma) \neq 0\}$ is finite.

We interpret $\mathcal{T}_\alpha$ as an $L_{s,\alpha}$-structure by defining

- $f \trianglelefteq g$ if and only if $f \subseteq g$. Write $f \perp g$ if $\neg(f \trianglelefteq g)$ and $\neg(g \trianglelefteq f)$.
- $f \wedge g = f|_{[\beta,\alpha)} = g|_{[\beta,\alpha)}$ where $\beta = \min\{\gamma : f|_{[\gamma,\alpha)} = g|_{[\gamma,\alpha)}\}$, if non-empty (note that β will not be a limit, by finite support). Define $f \wedge g$ to be the empty function if this set is empty (note that this cannot occur if α is a limit).
- $f <_{lex} g$ if and only if $f \triangleleft g$ or, $f \perp g$ with $\mathrm{dom}(f \wedge g) = [\gamma + 1, \alpha)$ and $f(\gamma) < g(\gamma)$
- For all $\beta < \alpha$, $P_\beta = \{f \in \mathcal{T}_\alpha : \mathrm{dom}(f) = [\beta, \alpha)\}$. P_α is only defined on $\mathcal{T}_\alpha$ if α is a successor, in which case it only contains the empty function.

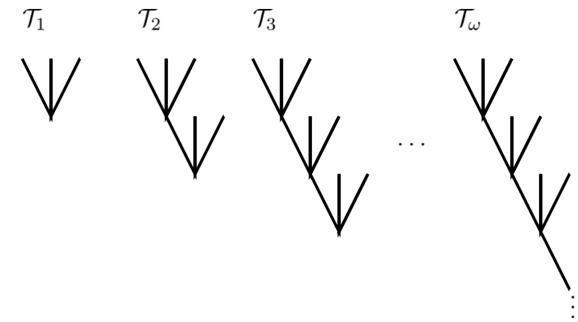

Figure 3.1: Illustration of the trees $\mathcal{T}_\alpha$

It is easy to check that for all $n < \omega$, $\mathcal{T}_n \cong \omega^{\leq n}$. For α infinite, however, $\mathcal{T}_\alpha$ will be ill-founded (as a partial order). In particular, P_0 names the level at the *top* of the tree, $P_{\beta+1}$ names the level immediately *below* P_β, and so on.

As many arguments in this paper will involve inductive constructions of trees of tuples indexed by $\mathcal{T}_\alpha$, it will be useful to fix notation as follows:

Definition 3.5.2. Suppose α is an ordinal.

1. (Restriction) If $w \subseteq \alpha$, the *restriction of $\mathcal{T}_\alpha$ to the set of levels w* is given by

$$\mathcal{T}_\alpha \restriction w = \{\eta \in \mathcal{T}_\alpha : \min(\mathrm{dom}(\eta)) \in w \text{ and } \beta \in \mathrm{dom}(\eta) \setminus w \implies \eta(\beta) = 0\}.$$

2. (Concatenation) If $\eta \in \mathcal{T}_\alpha$, $\mathrm{dom}(\eta) = [\beta + 1, \alpha)$, and $i < \omega$, let $\eta \frown \langle i \rangle$ denote the function $\eta \cup \{(\beta, i)\}$. We define $\langle i \rangle \frown \eta \in \mathcal{T}_{\alpha+1}$ to be $\eta \cup \{(\alpha, i)\}$. We write $\langle i \rangle$ for $\emptyset \frown \langle i \rangle$.
3. (Canonical inclusions) If $\alpha < \beta$, we define the map $\iota_{\alpha\beta} : \mathcal{T}_\alpha \to \mathcal{T}_\beta$ by $\iota_{\alpha\beta}(f) = f \cup \{(\gamma, 0) : \gamma \in \beta \setminus \alpha\}$.
4. (The all 0's path) If $\beta < \alpha$, then $\zeta_\beta \in \mathcal{T}_\alpha$ denotes the function with $\mathrm{dom}(\zeta_\beta) = [\beta, \alpha)$ and $\zeta_\beta(\gamma) = 0$ for all $\gamma \in [\beta, \alpha)$.

The function $i_{\alpha\beta}$ includes $\mathcal{T}_\alpha$ into $\mathcal{T}_\beta$ by adding zeros to the bottom of every node in $\mathcal{T}_\alpha$. Clearly if $\alpha < \beta < \gamma$, then $\iota_{\alpha\gamma} = \iota_{\beta\gamma} \circ \iota_{\alpha\beta}$. If β is a limit, then $\mathcal{T}_\beta$ is the direct limit of the $\mathcal{T}_\alpha$ for $\alpha < \beta$ along these maps. Visually, to get $\mathcal{T}_{\alpha+1}$ from $\mathcal{T}_\alpha$, one takes countably many copies of $\mathcal{T}_\alpha$ and adds a single root at the bottom.

Definition 3.5.3. Suppose I is an L'-structure, where L' is some language.

1. We say $(a_i : i \in I)$ is a set of *I-indexed indiscernibles* if whenever $(s_0, \ldots, s_{n-1})$, $(t_0, \ldots, t_{n-1})$ are tuples from I with

$$\mathrm{qftp}_{L'}(s_0, \ldots, s_{n-1}) = \mathrm{qftp}_{L'}(t_0, \ldots, t_{n-1}),$$

then we have

$$\mathrm{tp}(a_{s_0}, \ldots, a_{s_{n-1}}) = \mathrm{tp}(a_{t_0}, \ldots, a_{t_{n-1}}).$$

2. In the case that $L' = L_{s,\alpha}$ for some α, we say that an I-indexed indiscernible is *s-indiscernible*. As the only $L_{s,\alpha}$-structures we will consider will be trees, we will often refer to I-indexed indiscernibles in this case as *s-indiscernible trees*.

3. We say that I-indexed indiscernibles have the *modeling property* if, given any $(a_i : i \in I)$ from $\mathbb{M}$, there is an I-indexed indiscernible $(b_i : i \in I)$ in $\mathbb{M}$ *locally based* on $(a_i : i \in I)$ – i.e., given any finite set of formulas Δ from L and a finite tuple $(t_0, \dots, t_{n-1})$ from I, there is a tuple $(s_0, \dots, s_{n-1})$ from I so that

$$\mathrm{qftp}_{L'}(t_0, \dots, t_{n-1}) = \mathrm{qftp}_{L'}(s_0, \dots, s_{n-1})$$

and also

$$\mathrm{tp}_{\Delta}(b_{t_0}, \dots, b_{t_{n-1}}) = \mathrm{tp}_{\Delta}(a_{s_0}, \dots, a_{s_{n-1}}).$$

Fact 3.5.4. [KKS14, Theorem 4.3] Let denote I_s be the $L_{s,\omega}$-structure $(\omega^{<\omega}, \trianglelefteq, <_{lex}, \wedge, (P_\alpha)_{\alpha<\omega})$ with all symbols being given their intended interpretations and each P_α naming the elements of the tree at level α. Then I_s-indexed indiscernibles have the modeling property.

Remark 3.5.5. Note that the tree $\omega^{<\omega}$ is *not* the same tree as $\mathcal{T}_\omega$, which is ill-founded.

Corollary 3.5.6. *For any α, $\mathcal{T}_\alpha$-indexed indiscernibles have the modeling property.*

Proof. By Fact 3.5.4 and compactness. □

Definition 3.5.7. Suppose $(a_\eta)_{\eta \in \mathcal{T}_\alpha}$ is a tree of tuples, and C is a set of parameters.

1. We say $(a_\eta)_{\eta \in \mathcal{T}_\alpha}$ is *spread out over* C if for all $\eta \in \mathcal{T}_\alpha$ with $\mathrm{dom}(\eta) = [\beta+1, \alpha)$ for some $\beta < \alpha$, there is a global C-invariant type $q_\eta \supseteq \mathrm{tp}(a_{\trianglerighteq \eta \frown \langle 0 \rangle}/C)$ so that $(a_{\trianglerighteq \eta \frown \langle i \rangle})_{i<\omega}$ is a Morley sequence over C in q_η.

2. Suppose $(a_\eta)_{\eta \in \mathcal{T}_\alpha}$ is a tree which is spread out and s-indiscernible over C and for all $w, v \in [\alpha]^{<\omega}$ with $|w| = |v|$,

$$(a_\eta)_{\eta \in \mathcal{T}_\alpha \restriction w} \equiv_C (a_\eta)_{\eta \in \mathcal{T}_\alpha \restriction v}$$

then we say $(a_\eta)_{\eta \in \mathcal{T}_\alpha}$ is a *Morley tree* over C.

3. A *tree Morley sequence* over C is a C-indiscernible sequence of the form $(a_{\zeta_\beta})_{\beta<\alpha}$ for some Morley tree $(a_\eta)_{\eta \in \mathcal{T}_\alpha}$ over C.

Remark 3.5.8. If $(a_\eta)_{\eta \in \mathcal{T}_\alpha}$ is s-indiscernible over C, then, in order to be spread out over C, it suffices to have global C-invariant types as in (1) for all η identically zero—i.e. those nodes in the tree of the form ζ_β for some $\beta < \alpha$. Note that the condition in (2) forces $(a_{\zeta_\beta})_{\beta<\alpha}$ to be C-indiscernible—in fact, (1) and (2) together can be shown to be equivalent to demanding that the tree is indiscernible with respect to the language $L = \langle \trianglelefteq, <_{lex}, \wedge, \leq_{len} \rangle$, where $\leq_{len}$ is interpreted as the pre-order which compares the lengths of nodes in the tree. Finally, in (3) we speak of $(a_{\zeta_\beta})_{\beta<\alpha}$, the sequence indexed by the all-zeroes path in the tree, simply because this is a convenient choice of a path. In an s-indiscernible tree over C, any two paths will have the same type over C. Hence, (3) may be stated more succinctly as: a tree Morley sequence over C is a path in some Morley tree over C.

Lemma 3.5.9. *Suppose $(a_i)_{i<\omega}$ is a tree Morley sequence over C.*

1. *If $a_i = (b_i, c_i)$ for all $i < \omega$, where the b_i's are all initial subtuples of a_i of the same length, then $(b_i)_{i<\omega}$ is a tree Morley sequence over C.*
2. *Given $1 \leq n < \omega$, suppose $d_i = (a_{n\cdot i}, a_{n\cdot i+1}, \dots, a_{n\cdot i+n-1})$. Then $(d_i)_{i<\omega}$ is a tree Morley sequence over C.*

Proof. (1) is immediate from the definition: s-indiscernibility, spread-outness, and being a Morley tree over C are all preserved under taking subtuples.

(2) Suppose $(a_\eta)_{\eta\in\mathcal{T}_\omega}$ is a Morley tree over C with $a_{\zeta_i} = a_i$. Define a function $j : \mathcal{T}_\omega \to \mathcal{T}_\omega$ so that if $\eta \in \mathcal{T}_\omega$ with $\mathrm{dom}(\eta) = [k, \omega)$, then $\mathrm{dom}(j(\eta)) = [n(k+1), \omega)$ and

$$j(\eta)(l) = \begin{cases} \eta\left(\frac{l}{n} - 1\right) & \text{if } n|l \\ 0 & \text{otherwise} \end{cases}$$

for all $l \in [n(k+1), \omega)$. Define $(b_\eta)_{\eta\in\mathcal{T}_\omega}$ by

$$b_\eta = (a_{j(\eta)}, a_{j(\eta)\frown\langle 0\rangle}, \dots, a_{j(\eta)\frown 0^{n-1}}).$$

It is easy to check that this is also an s-indiscernible tree over M (more formally, this construction corresponds to the n-fold elongation of the tree $(a_\eta)_{\eta\in\mathcal{T}_\omega}$ as defined in Chapter 1 so $(b_\eta)_{\eta\in\mathcal{T}_\omega}$ is s-indiscernible over M by Proposition I.2.1(1) there). It is also easy to check that $(b_\eta)_{\eta\in\mathcal{T}_\omega}$ is spread out over M. Finally, the tree $(b_\eta)_{\eta\in\mathcal{T}_\omega}$ is also a Morley tree over M: given $w \in [\omega]^{<\omega}$, let $w' = \{n(k+1) - l : k \in w, l < n\}$. Then if $w, v \in [\omega]^{<\omega}$ and $|w| = |v|$, then $|w'| = |v'|$ so $(a_\eta)_{\eta\in\mathcal{T}_\omega\restriction w'} \equiv_C (a_\eta)_{\eta\in\mathcal{T}_\omega\restriction v'}$ so $(b_\eta)_{\eta\in\mathcal{T}_\omega\restriction w} \equiv_C (b_\eta)_{\eta\in\mathcal{T}_\omega\restriction v}$. It follows that $(b_{\zeta_i})_{i<\omega}$ is a tree Morley sequence over C. We have

$$\begin{aligned} b_{\zeta_i} &= (a_{\zeta_{n(i+1)}}, a_{\zeta_{n(i+1)}\frown\rangle 0\rangle}, \dots, a_{\zeta_{n(i+1)}\frown 0^{n-1}}) \\ &= (a_{n(i+1)}, a_{n(i+1)-1}, \dots, a_{n(i+1)-(n-1)}), \end{aligned}$$

so by reversing the order of the tuple, we deduce that $(d_i)_{i<\omega}$ is a tree Morley sequence over M. □

From the existence of a sufficiently large tree which is spread out and s-indiscernible over M, one can obtain a Morley tree which is based on it. The proof is via a standard Erdős-Rado argument. We follow the argument of [GIL02, Theorem 1.13].

Lemma 3.5.10. *Suppose $(a_\eta)_{\eta\in\mathcal{T}_\kappa}$ is a tree of tuples, spread out and s-indiscernible over M. If κ is sufficiently large, then there is a Morley tree $(b_\eta)_{\eta\in\mathcal{T}_\omega}$ so that for all $w \in [\omega]^{<\omega}$, there is $v \in [\kappa]^{<\omega}$ so that*

$$(a_\eta)_{\eta\in\mathcal{T}_\kappa\restriction v} \equiv_M (b_\eta)_{\eta\in\mathcal{T}_\omega\restriction w}.$$

Proof. Let $\lambda = 2^{|M|+|T|}$ and set $\kappa = \beth_{\lambda^+}(\lambda)$. Given a tree $(a_\eta)_{\eta \in \mathcal{T}_\kappa}$ s-indiscernible and spread out over M, let

$$\Gamma_n = \{\text{tp}((a_\eta)_{\eta \in \mathcal{T}_\omega \restriction w}/M) : w \in [\kappa]^n\}.$$

By induction on n, we will find a sequence of types $p_n \in \Gamma_n$ so that

$$\Delta(x_\eta : \eta \in \mathcal{T}_\omega) = \bigcup_{n<\omega} \bigcup_{w \in [\omega]^n} p_n(x_\eta : \eta \in \mathcal{T}_\omega \restriction w)$$

is consistent. Construct by induction on n cofinal subsets $F_n \subseteq \lambda^+$ and subsets $X_{\xi,n} \subseteq \kappa$ so that

1. $F_{n+1} \subseteq F_n$.
2. $|X_{\xi,n}| > \beth_\alpha(\lambda)$ when ξ is the αth element of F_n.
3. If $w \in [X_{\xi,n}]^n$, then $(a_\eta)_{\eta \in \mathcal{T}_\kappa \restriction w} \models p_n$.
4. $|F_n| = \lambda^+$.

For $n = 0$, we let $F_0 = \lambda^+$ and $X_{\xi,0} = \kappa$ for all $\xi < \lambda^+$. Suppose F_n and $(X_{\xi,n})_{\xi \in F_n}$ have been constructed. Write $F_n = \{\xi_\alpha : \alpha < \lambda^+\}$ where the ξ_α enumerate F_n in increasing order. Then for all $\alpha < \lambda^+$,

$$|X_{\xi_{\alpha+n+1},n}| > \beth_{\alpha+n+1}(\lambda).$$

For a moment, fix $\xi = \xi_{\alpha+n+1}$. Define a coloring on $[X_{\xi,n}]^{n+1}$ by

$$w \mapsto \text{tp}((a_\eta)_{\eta \in \mathcal{T}_\kappa \restriction w}/M).$$

This is a coloring with at most λ many colors so by Erdős-Rado there is a homogeneous subset $X_{\xi,n+1} \subseteq X_{\xi,n}$ with $|X_{\xi,n+1}| > \beth_\alpha(\lambda)$. Let $p_{n+1,\alpha+n+1}$ denote its constant value. By the pigeonhole principle, as the set of possible values is λ and $\{\alpha + n + 1 : \alpha < \lambda^+\}$ has size λ^+, there must be some subset $Y \subseteq \{\alpha+n+1 : \alpha < \lambda^+\}$ of cardinality λ^+ so that $\beta, \beta' \in Y$ implies $p_{n+1,\beta} = p_{n+1,\beta'}$. Let $p_{n+1} = p_{n+1,\beta}$ for some/all $\beta \in Y$. Put $F_{n+1} = \{\xi_\beta : \beta \in Y\}$. Then p_{n+1}, F_{n+1}, and $(X_{\xi,n+1})_{\xi \in F_{n+1}}$ clearly satisfy the requirements.

By compactness, this shows that $\Delta(x_\eta : \eta \in \mathcal{T}_\omega)$ is consistent. Let $(b_\eta)_{\eta \in \mathcal{T}_\omega}$ be a realization—now to show $(b_\eta)_{\eta \in \mathcal{T}_\omega}$ is a Morley tree over M, we must show that $(b_\eta)_{\eta \in \mathcal{T}_\omega}$ is s-indiscernible and spread out over M. To see that it is spread out over M, fix any $\eta \in \mathcal{T}_\omega$ with $\text{dom}(\eta) = [n+1, \omega)$. Setting $w = \{0, \ldots, n\}$, there is $v \in [\kappa]^{<\omega}$, $v = \{\alpha_0 < \ldots < \alpha_n\}$ so that $(b_\nu)_{\nu \in \mathcal{T}_\omega \restriction w} \equiv_M (a_\nu)_{\nu \in \mathcal{T}_\kappa \restriction v}$. If $\nu_i \in \mathcal{T}_\kappa$ has domain $[\alpha_0, \kappa)$, $\nu_i(\alpha_0) = i$ and ν_i is identically zero elsewhere, $\langle a_{\unrhd \nu_i} : i < \omega \rangle$ is a Morley sequence over M in an M-invariant type. It follows that $\langle b_{\unrhd \eta \frown \langle i \rangle} : i < \omega \rangle$ is also a Morley sequence in an M-invariant type, which establishes spread-outness of the tree. Checking that the tree is s-indiscernible over M is entirely similar. $\square$

The symmetry characterization of NSOP_1

In this subsection, we prove a version of Kim's lemma for tree Morley sequences and use it to prove that Kim-independence is symmetric over models in an NSOP_1 theory. Lemma 3.5.11 is the key step, showing that tree Morley sequences exist under certain assumptions. The method of proof is an inductive construction of a spread out s-indiscernible tree, from which a Morley tree (and hence a tree Morley sequence) can then be extracted. This basic proof-strategy will be repeated several times throughout the paper.

Lemma 3.5.11. *Suppose T is $NSOP_1$, $M \models T$, and $a \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^K_M b$. For any ordinal $\alpha \geq 1$, there is a spread out s-indiscernible tree $(c_\eta)_{\eta \in \mathcal{T}_\alpha}$ over M, so that if $\eta \lhd \nu$ and $dom(\nu) = \alpha$, then $c_\eta c_\nu \equiv_M ab$.*

Proof. We will argue by induction on α. For the case $\alpha = 1$, fix $q \supseteq \mathrm{tp}(b/M)$, a global M-invariant type. Let $\langle b_i : i < \omega \rangle \models q^{\otimes \omega}|_M$. As $a \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^K_M b$, we may assume this sequence is Ma-indiscernible. Put $c^1_\emptyset = a$ and $c^1_{\langle i \rangle} = b_i$. It is now easy to check that $(c^1_\eta)_{\eta \in \mathcal{T}_1}$ is a spread out s-indiscernible tree satisfying the requirements.

Suppose for α we have constructed $(c^\beta_\eta)_{\eta \in \mathcal{T}_\beta}$ for $1 \leq \beta \leq \alpha$ such that, if $\gamma < \beta \leq \alpha$ and $\eta \in \mathcal{T}_\gamma$ then $c^\gamma_\eta = c^\beta_{\iota_{\gamma\beta}(\eta)}$. By spread-outness, we know that $\langle c^\alpha_{\unrhd \langle i \rangle} : i < \omega \rangle$ is an M-invariant Morley sequence which is, by s-indiscernibility over M, $Mc^\alpha_\emptyset$-indiscernible. Therefore, $c^\alpha_\emptyset \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^K_M (c^\alpha_{\unrhd \langle i \rangle})_{i<\omega}$. By extension (Proposition 6.3.18), we may find some $c' \equiv_{M(c^\alpha_{\unrhd\langle i \rangle})_{i<\omega}} c^\alpha_\emptyset$ so that

$$c' \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^K_M (c^\alpha_\eta)_{\eta \in \mathcal{T}_\alpha}.$$

Choose a global M-invariant type $q \supseteq \mathrm{tp}((c^\alpha_\eta)_{\eta \in \mathcal{T}_\alpha}/M)$. Let $\langle (c^\alpha_{\eta,i})_{\eta \in \mathcal{T}_\alpha} : i < \omega \rangle \models q^{\otimes \omega}|_M$ with $c^\alpha_{\eta,0} = c^\alpha_\eta$ for all $\eta \in \mathcal{T}_\alpha$. By the chain condition (Lemma 3.3.21), we can find $c'' \equiv_{M(c^\alpha_\eta)_{\eta \in \mathcal{T}_\alpha}} c'$ so that $c'' \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}^K_M (c^\alpha_{\eta,i})_{\eta \in \mathcal{T}_\alpha, i<\omega}$ and $\langle (c^\alpha_{\eta,i})_{\eta \in \mathcal{T}_\alpha} : i < \omega \rangle$ is Mc''-indiscernible. Define a new tree $(d_\eta)_{\eta \in \mathcal{T}_{\alpha+1}}$ by setting $d_\emptyset = c''$ and $d_{\iota_{\alpha\alpha+1}(\eta)} = c^\alpha_\eta$ for all $\eta \in \mathcal{T}_\alpha$. Then let $(c^{\alpha+1}_\eta)_{\eta \in \mathcal{T}_{\alpha+1}}$ be a tree s-indiscernible over M locally based on $(d_\eta)_{\eta \in \mathcal{T}_\alpha}$. By an automorphism, we may assume $c^{\alpha+1}_{\iota_{\alpha\alpha+1}(\eta)} = c^\alpha_\eta$ for all $\eta \in \mathcal{T}_\alpha$. This satisfies our requirements.

Finally, suppose for δ limit we have constructed $(c^\beta_\eta)_{\eta \in \mathcal{T}_\beta}$ for $1 \leq \beta < \delta$ such that, if $\gamma < \beta < \delta$ and $\eta \in \mathcal{T}_\gamma$ then $c^\gamma_\eta = c^\beta_{\iota_{\gamma\beta}(\eta)}$. If $\eta \in \mathcal{T}_\delta$, then for some $\beta < \delta$, there is $\nu \in \mathcal{T}_\beta$ so that $\iota_{\beta\delta}(\nu) = \eta$. Then put $c^\delta_\eta = c^\beta_\nu$. This defines for all $\beta \leq \delta$ an s-indiscernible tree $(c^\beta_\eta)_{\eta \in \mathcal{T}_\eta}$ satisfying our requirements. □

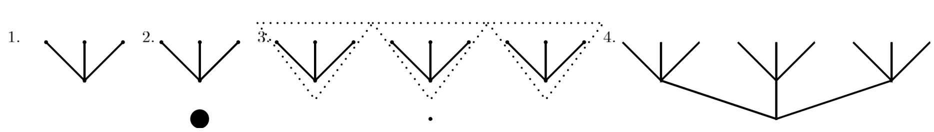

Figure 3.2: The construction of the tree indexed by $\mathcal{T}_2$ in stages.

The figure displays the construction of the tree indexed by $\mathcal{T}_2$ in stages: 1. The tree indexed by $\mathcal{T}_1$. 2. Using extension to obtain a new base point. 3. Taking a Morley sequence in the given tree, indiscernible over the new base point. 4. Extracting an s-indiscernible tree to obtain a spread out, s-indiscernible tree indexed by $\mathcal{T}_2$.

Lemma 3.5.12. *Suppose T is NSOP$_1$, $M \models T$, and $a \mathop{\smile\!\!\!\!\!|}^K_M b$. Then there is a tree Morley sequence $(a_i)_{i<\omega}$ which is Mb-indiscernible with $a_0 = a$.*

Proof. By Lemma 3.5.11, for arbitrarily large cardinals κ, there is a tree $(c_\eta)_{\eta \in \mathcal{T}_\kappa}$ which is spread out and s-indiscernible over M so that if $\eta \lhd \nu$ and $\mathrm{dom}(\nu) = \kappa$ then $c_\eta c_\nu \equiv_M ab$. Note that $\mathcal{T}' = \mathcal{T}_\kappa \setminus \{\nu \in \mathcal{T}_\kappa : \mathrm{dom}(\nu) = \kappa\} = \{\eta \in \mathcal{T}_\kappa : \mathrm{dom}(\eta) \subseteq [1,\kappa)\}$ is isomorphic to $\mathcal{T}_\kappa$. So we may enumerate $(c_\eta)_{\eta \in \mathcal{T}'}$ as $(d_\eta)_{\eta \in \mathcal{T}_\kappa}$. Note that for all $\eta \in \mathcal{T}_\kappa$, $d_\eta \equiv_M a$ and $d_{\zeta_\alpha} = c_{\zeta_{1+\alpha}}$ for all $\alpha < \kappa$. By Lemma 3.5.10, there is a Morley tree over M $(d'_\eta)_{\eta \in \mathcal{T}_\omega}$ so that for all $w \in [\omega]^{<\omega}$ there is $v \in [\kappa]^{<\omega}$ so that $(d_\eta)_{\eta \in \mathcal{T}_\kappa \restriction v} \equiv_M (d'_\eta)_{\eta \in \mathcal{T}_\omega \restriction w}$.

Let $p(x;a) = \mathrm{tp}(b/Ma)$. We claim $\bigcup_{i<\omega} p(x;d'_{\zeta_i})$ is consistent. Given n, let $w = \{0,\ldots,n-1\}$. Find $v \in [\kappa]^{<\omega}$ so that $(d_\eta)_{\eta \in \mathcal{T}_\kappa \restriction v} \equiv_M (d'_\eta)_{\eta \in \mathcal{T}_\omega \restriction w}$. If $v = \{\alpha_0,\ldots,\alpha_{n-1}\}$, then for $i < n$ we have $d_{\zeta_{\alpha_i}} = c_{1+\zeta_{\alpha_i}}$. Then because $c_{\zeta_{1+\alpha_i}} c_{\zeta_0} \equiv_M ab$ for all $i < n$, we have $c_{\zeta_0} \models \bigcup_{i<n} p(x;d_{\zeta_{\alpha_i}})$. This shows $\bigcup_{i<n} p(x;d_{\zeta_{\alpha_i}})$ is consistent and hence $\bigcup_{i<n} p(x;d'_{\zeta_i})$ is consistent. The claim follows by compactness.

Let $b' \models \bigcup_{i<\omega} p(x;d'_{\zeta_i})$. Extract from $(d'_{\zeta_i})_{i<\omega}$ an Mb'-indiscernible sequence $(a_i)_{i<\omega}$. As $(a_i)_{i<\omega} \equiv_M (d'_{\zeta_i})_{i<\omega}$, we know $(a_i)_{i<\omega}$ is a tree Morley sequence. By an automorphism, we may assume $b' = b$ and $a_0 = a$. □

Proposition 3.5.13. *Suppose T is NSOP$_1$ and $M \models T$. Suppose $(a_i)_{i<\omega}$ is a tree Morley sequence over M. Then $\{\varphi(x;a_i) : i < \omega\}$ is inconsistent if and only if $\varphi(x;a_0)$ Kim-divides over M.*

Proof. Suppose $(a_i)_{i<\omega}$ is a tree Morley sequence over M. Let $(a_\eta)_{\eta \in \mathcal{T}_\omega}$ be a Morley tree over M with $a_{\zeta_i} = a_i$. Let $\eta_i \in \mathcal{T}_\omega$ be the function with $\mathrm{dom}(\eta_i) = [i,\omega)$ and

$$\eta_i(j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Consider the sequence $I = (a_{\eta_i}, a_{\zeta_i})_{i<\omega}$. Because $(a_\eta)_{\eta \in \mathcal{T}_\omega}$ is a Morley tree over M, I is an M-indiscernible sequence. Moreover, by s-indiscernibility, $a_{\eta_0} \equiv_{MI_{>0}} a_{\zeta_0}$. By indiscernibility, for all i, we have $a_{\eta_i} \equiv_{MI_{>i}} a_{\zeta_i}$. By NSOP$_1$, it follows that $\{\varphi(x;a_{\eta_i}) : i < \omega\}$ is consistent if and only if $\{\varphi(x;a_{\zeta_i}) : i < \omega\}$ is consistent: if exactly one of them is consistent, then we have SOP$_1$ by Proposition 5.2.1.

Because $(a_\eta)_{\eta \in \mathcal{T}_\omega}$ is a spread out tree over M, $a_{\eta_i} \mathop{\smile\!\!\!\!\!|}^i_M a_{\eta_{<i}}$ for all i. Using the fact that $(a_{\eta_i})_{i<\omega}$ is an M-indiscernible sequence and the compactness of the space of M-invariant types, we have $(a_{\eta_i})_{i<\omega}$ is a Morley sequence in some global M-invariant type extending $\mathrm{tp}(a/M)$, so $\varphi(x;a)$ Kim-divides over M if and only if $\{\varphi(x;a_{\eta_i}) : i < \omega\}$ is inconsistent. □

Corollary 3.5.14. *(Kim's lemma for tree Morley sequences) Suppose T is NSOP$_1$ and $M \models T$. The following are equivalent:*

1. *$\varphi(x;a)$ Kim-divides over M.*

2. *For some tree Morley sequence $(a_i)_{i<\omega}$ over M with $a_0 = a$, $\{\varphi(x;a_i) : i < \omega\}$ is inconsistent.*

3. *For every tree Morley sequence $(a_i)_{i<\omega}$ over M with $a_0 = a$, $\{\varphi(x;a_i) : i < \omega\}$ is inconsistent.*

Corollary 3.5.15. *(Chain condition for tree Morley sequences) Suppose T is NSOP$_1$ and $M \models T$. If $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ and $I = (b_i)_{i<\omega}$ is a tree Morley sequence over M with $b_0 = b$, then there is $a' \equiv_{Mb} a$ so that $a' \mathop{\smile\!\!\!\!\!\!\mid}^K_M I$ and I is Ma'-indiscernible.*

Proof. The proof is identical to Proposition 3.3.21 above, since, by Lemma 3.5.9, $\langle (b_{kn+n-1}, b_{kn+n-2}, \ldots, b_{kn}) : k < \omega\rangle$ is a tree Morley sequence over M. □

Theorem 3.5.16. *(Symmetry) Suppose T is a complete theory. The following are equivalent:*

1. *T is NSOP$_1$.*

2. *$\mathop{\smile\!\!\!\!\!\!\mid}^K$ is symmetric over models: for any $M \models T$ and tuples a, b from $\mathbb{M}$, $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b \iff b \mathop{\smile\!\!\!\!\!\!\mid}^K_M a$.*

3. *$\mathop{\smile\!\!\!\!\!\!\mid}^K$ enjoys the following weak symmetry property: for any $M \models T$ and tuples a, b from $\mathbb{M}$, $a \mathop{\smile\!\!\!\!\!\!\mid}^i_M b$ implies $b \mathop{\smile\!\!\!\!\!\!\mid}^K_M a$.*

Proof. (1) $\iff$ (3) is Proposition 3.3.22 and (2) $\implies$ (3) is immediate from the fact that $a \mathop{\smile\!\!\!\!\!\!\mid}^i_M b$ implies $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$.

(1) $\implies$ (2). Suppose T is NSOP$_1$. Assume towards contradiction that $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ and $b \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M a$. By Lemma 3.5.12, there is a tree Morley sequence over M with $a_0 = a$ which is Mb-indiscernible. Since $b \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M a$, there is some $\varphi(x;a) \in \mathrm{tp}(b/Ma)$ which Kim-divides over M. By Corollary 3.5.14, $\{\varphi(x;a_i) : i < \omega\}$ is inconsistent. But $\models \varphi(b;a_i)$ for all $i < \omega$ by indiscernibility, a contradiction. □

Corollary 3.5.17. *Assume the complete theory T is NSOP$_1$ and $M \models T$. Then*

$$a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b \iff acl(a) \mathop{\smile\!\!\!\!\!\!\mid}^K_M b \iff a \mathop{\smile\!\!\!\!\!\!\mid}^K_M acl(b).$$

Proof. By symmetry, it is enough to prove $\mathrm{acl}(a) \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$, assuming $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$. If $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$, there is a Morley sequence in an M-invariant type $\langle b_i : i < \omega\rangle$ with $b_0 = b$ which is Ma-indiscernible. Then it is automatically $M\mathrm{acl}(a)$-indiscernible so $\mathrm{acl}(a) \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$. □

3.6 The independence theorem

The full independence theorem will be deduced from a weak independence theorem, which has an easy proof:

Proposition 3.6.1. *Assume T is NSOP$_1$. Then $\mathop{\smile\!\!\!\!\!\!\mid}^K$ satisfies the following weak independence theorem over models: if $M \models T$, $a \equiv_M a'$, $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$, $a' \mathop{\smile\!\!\!\!\!\!\mid}^K_M c$ and $b \mathop{\smile\!\!\!\!\!\!\mid}^u_M c$, then there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\!\!\!\!\!\!\mid}^K_M bc$.*

Proof. Suppose T is NSOP$_1$ and fix $M \models T$ and tuples a, a', b, c so that $a \equiv_M a'$, $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$, $a' \mathop{\smile\!\!\!\!\!\!\mid}^K_M c$ and $b \mathop{\smile\!\!\!\!\!\!\mid}^u_M c$.

Claim: There is c' so that $ac' \equiv_M a'c$ and $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M bc'$.

Proof of claim: By symmetry, it suffices to find c' with $ac' \equiv_M a'c$ and $bc' \mathop{\smile\!\!\!\!\!\!\mid}^K_M a$. Let $p(x; a') = \text{tp}(c/Ma')$. By invariance, we know $p(x; a)$ does not Kim-fork over M. We have to show

$$p(x; a) \cup \{\neg\varphi(x, b; a) : \varphi(x, y; a) \in L(Ma) \text{ Kim-divides over } M\}$$

is consistent. If not, then by compactness and Kim-forking = Kim-dividing, we must have

$$p(x; a) \vdash \varphi(x, b; a),$$

for some φ where $\varphi(x, y; a)$ Kim-divides over M. By symmetry, $b \mathop{\smile\!\!\!\!\!\!\mid}^K_M a$, so there is some M-invariant Morley sequence $(a_i)_{i<\omega}$ with $a_0 = a$ which is moreover Mb-indiscernible. Then we have

$$\bigcup_{i<\omega} p(x; a_i) \vdash \{\varphi(x, b; a_i) : i < \omega\}.$$

As $p(x; a)$ does not Kim-fork over M, we know $\bigcup_{i<\omega} p(x; a_i)$ is consistent. But, by Kim's lemma for Kim-dividing, we know $\{\varphi(x, y; a_i) : i < \omega\}$ is inconsistent and *a fortiori* $\{\varphi(x, b; a_i) : i < \omega\}$ is inconsistent, a contradiction. So the given partial type is consistent. Let c' realize it. Then $ac' \equiv_M a'c$ and $c'b \mathop{\smile\!\!\!\!\!\!\mid}^K_M a$, which proves the claim. □

As $b \mathop{\smile\!\!\!\!\!\!\mid}^u_M c$, by left extension, there is $c'' \equiv_{Mb} c$ with $bc' \mathop{\smile\!\!\!\!\!\!\mid}^u_M c''$. Then by right extension and automorphism, we can choose some b'' so that $bc' \equiv_M b''c''$ and $bc' \mathop{\smile\!\!\!\!\!\!\mid}^u_M b''c''$. As $bc' \mathop{\smile\!\!\!\!\!\!\mid}^u_M b''c''$ and $bc' \equiv_M b''c''$, it follows that $(b''c'', bc')$ starts a Morley sequence I in some global M-finitely satisfiable (hence M-invariant) type. As $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M bc'$, we may, by the chain condition (Proposition 3.3.21) find some $a_* \equiv_{Mbc'} a$ so that I is Ma_*-indiscernible and $a_* \mathop{\smile\!\!\!\!\!\!\mid}^K_M I$. Then, we obtain $a_* \equiv_{Mb} a$, $a_*c'' \equiv_M a'c$, and $a_* \mathop{\smile\!\!\!\!\!\!\mid}^K_M bc''$. By construction, $c'' \equiv_{Mb} c$ so there is $\sigma \in \text{Aut}(\mathbb{M}/Mb)$ with $\sigma(c'') = c$. Then $\sigma(a_*) \mathop{\smile\!\!\!\!\!\!\mid}^K_M bc$, $\sigma(a_*) \equiv_{Mb} a$, and $\sigma(a_*) \equiv_{Mc} a'$, which shows that the weak independence theorem over models holds for T. □

Lemma 3.6.2. *Suppose T is NSOP$_1$, $M \models T$, and $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$. Fix an ordinal α and any $q \supseteq tp(b/M)$, a global M-invariant type. If $(b_\eta)_{\eta \in \mathcal{T}_\alpha}$ is a tree, spread out over M, so that,* *for all $\nu \in \mathcal{T}_\alpha$, $b_\nu \models q|_{Mb_{\rhd\nu}}$, then, writing $p(x;b)$ for $tp(a/Mb)$, we have*

$$\bigcup_{\eta\in\mathcal{T}_\alpha} p(x;b_\eta)$$

is consistent and non-Kim-forking over M.

Proof. The proof is by induction on α. For $\alpha = 0$, there is nothing to show. For α limit, it follows by induction, using that $\mathcal{T}_\alpha$ is the direct limit of the $\mathcal{T}_\beta$ for $\beta < \alpha$ along the maps $\iota_{\beta\alpha}$. Now suppose given $(b_\eta)_{\eta\in\mathcal{T}_{\alpha+1}}$ as in the statement. We know that $b_{\unrhd\zeta_\alpha} = (b_{\iota_{\alpha\alpha+1}(\eta)})_{\eta\in\mathcal{T}_\alpha}$ is a tree spread out over M so that, for all $\nu \in \mathcal{T}_\alpha$, $b_{\iota_{\alpha\alpha+1}(\nu)} \models q|_{M(b_{\iota_{\alpha\alpha+1}(\eta)})_{\eta\unrhd\nu}}$. Note that $\emptyset \frown \langle 0\rangle = \zeta_\alpha$. By induction, then,

$$\bigcup_{\nu\unrhd\emptyset\frown\langle 0\rangle} p(x;b_\nu)$$

is consistent and non-Kim-forking over M. By spread outness over M, $\langle b_{\unrhd\emptyset\frown\langle i\rangle} : i<\omega\rangle$ is a Morley sequence in some global M-invariant type. By the chain condition,

$$\bigcup_{i<\omega}\bigcup_{\nu\unrhd\emptyset\frown\langle i\rangle} p(x;b_\nu)$$

is consistent and non-Kim-forking over M. As $b_\emptyset \models q|_{M(b_{\nu\rhd\emptyset})}$, it follows by Proposition 3.6.1 that

$$p(x;b_\emptyset)\cup\bigcup_{i<\omega}\bigcup_{\nu\unrhd\emptyset\frown\langle i\rangle} p(x;b_\nu)$$

is consistent and non-Kim-forking over M. Unwinding definitions, this says

$$\bigcup_{\eta\in\mathcal{T}_{\alpha+1}} p(x;b_\eta)$$

is consistent and non-Kim-forking over M, completing the proof. □

Remark 3.6.3. In the above proof, the hypothesis that $b_\nu \models q|_{Mb_{\rhd\nu}}$ is used to apply the weak independence theorem (Proposition 3.6.1). Once one has proved the full independence theorem (Theorem 3.6.5), the same proof gives $\bigcup_{\eta\in\mathcal{T}_\alpha} p(x;b_\eta)$ is consistent and non-Kim-forking over M, just under the hypothesis that $(b_\eta)_{\eta\in\mathcal{T}_\alpha}$ is $s-$indiscernible and spread out over M, since $b_\nu \mathop{\smile\!\!\!\!\!\mid}^K_M b_{\rhd\nu}$ in any tree s-indiscernible and spread out over M.

Lemma 3.6.4. *(Zig-zag Lemma) Suppose the complete theory T is NSOP$_1$, $M \models T$ and $b \mathop{\smile\!\!\!\!\!\mid}^K_M b'$. Then for any global M-invariant type $q \supseteq tp(b/M)$, there is a tree Morley sequence over M $(b_i, b'_i)_{i<\omega}$ starting with (b,b') so that*

1. *If $i \le j$, then $b_i b'_j \equiv_M bb'$.*

2. *If $i > j$, then $b_i \models q|_{Mb'_j}$.*

Proof. Fix $q \supseteq \mathrm{tp}(b/M)$ and let $p(x;b) = \mathrm{tp}(b'/Mb)$. By recursion on α, we will construct trees $(c^\alpha_\eta, d^\alpha_\eta)_{\eta \in \mathcal{T}_\alpha}$ so that, for all α

1. If $\eta \in \mathcal{T}_\alpha$, then
$$c^\alpha_\eta \models q|_{Mc^\alpha_{\rhd \eta} d^\alpha_{\rhd \eta}}$$

2. If $\eta \in \mathcal{T}_\alpha$, then
$$d^\alpha_\eta \models \bigcup_{\nu \unrhd \eta} p(x; c^\alpha_\nu)$$

3. $(c^\alpha_\eta, d^\alpha_\eta)_{\eta \in \mathcal{T}_\alpha}$ is spread out and s-indiscernible over M

4. If $\beta < \alpha$ then $(c^\alpha_{\iota_{\beta\alpha}(\eta)}, d^\alpha_{\iota_{\beta\alpha}(\eta)}) = (c^\beta_\eta, d^\beta_\eta)$ for all $\eta \in \mathcal{T}_\beta$.

To start, define $(c^0_\emptyset, d^0_\emptyset) = (b, b')$. This defines $(c^0_\eta, d^0_\eta)_{\eta \in \mathcal{T}_0}$.

Now suppose given $(c^\alpha_\eta, d^\alpha_\eta)_{\eta \in \mathcal{T}_\alpha}$. Let $\langle (c^\alpha_{\eta,i}, d^\alpha_{\eta,i}) : i < \omega \rangle$ be an M-invariant Morley sequence with $(c^\alpha_{\eta,0}, d^\alpha_{\eta,0})_{\eta \in \mathcal{T}_\alpha} = (c^\alpha_\eta, d^\alpha_\eta)_{\eta \in \mathcal{T}_\alpha}$. Pick c_* so that
$$c_* \models q|_{M(c^\alpha_{\eta,i}, d^\alpha_{\eta,i})_{\eta \in \mathcal{T}_\alpha, i<\omega}}.$$
Then, by Lemma 3.6.2, we may choose d_* so that
$$d_* \models \bigcup_{\substack{\eta \in \mathcal{T}_\alpha \\ i<\omega}} p(x; c^\alpha_{\eta,i}) \cup p(x; c_*).$$
Define a tree $(e_\eta, f_\eta)_{\eta \in \mathcal{T}_{\alpha+1}}$ by
$$\begin{aligned} (e_\emptyset, f_\emptyset) &= (c_*, d_*) \\ (e_{\langle i \rangle \frown \eta}, f_{\langle i \rangle \frown \eta}) &= (c^\alpha_{\eta,i}, d^\alpha_{\eta,i}). \end{aligned}$$
Finally, let $(c^{\alpha+1}_\eta, d^{\alpha+1}_\eta)_{\eta \in \mathcal{T}_{\alpha+1}}$ be a tree s-indiscernible over M locally based on this tree. By an automorphism, we may assume that $c^{\alpha+1}_{\iota_{\alpha\alpha+1}(\eta)} = c^\alpha_\eta$ for all $\eta \in \mathcal{T}_\alpha$. This satisfies the requirements.

Finally, arriving to stage δ for δ limit, we simply define $(c^\delta_\eta, d^\delta_\eta)_{\eta \in \mathcal{T}_\delta}$ by stipulating $(c^\delta_{\iota_{\beta\delta}(\eta)}, d^\delta_{\iota_{\beta\delta}(\eta)}) = (c^\beta_\eta, d^\beta_\eta)$ for all $\beta < \delta$. By the coherence condition (4), this is well-defined, and satisfies the requirements. We conclude by extracting a Morley tree, by Lemma 3.5.10. □

Theorem 3.6.5. *Suppose T is a complete theory. The following are equivalent:*

1. *T is NSOP$_1$.*

2. *$\mathop{\smile\!\!\!\!\!\!|}^K$ satisfies the independence theorem over models: if $M \models T$, $a \equiv_M a'$, $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$, $a' \mathop{\smile\!\!\!\!\!\!|}^K_M c$, and $b \mathop{\smile\!\!\!\!\!\!|}^K_M c$, then there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\!\!\!\!\!\!|}^K_M bc$.*

Proof. (2) $\implies$ (1) follows from Theorem I.5.1, using that $\mathop{\smile\!\!\!\!\!\mid}^i$ implies $\mathop{\smile\!\!\!\!\!\mid}^K$.

(1) $\implies$ (2): Assume T is NSOP$_1$. Suppose $M \models T$, $a \equiv_M a'$, and $a \mathop{\smile\!\!\!\!\!\mid}^K_M b$, $a' \mathop{\smile\!\!\!\!\!\mid}^K_M c$ and $b \mathop{\smile\!\!\!\!\!\mid}^K_M c$. We must show there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\!\!\!\!\!\mid}^K_M bc$. Let $p_0(x; b) = \text{tp}(a/Mb)$ and $p_1(x; c) = \text{tp}(a'/Mc)$. Suppose towards contradiction that $p_0(x; b) \cup p_1(x; c)$ Kim-forks over M. Let $q \supseteq \text{tp}(b/M)$ be a global type finitely satisfiable in M. In particular, q is M-invariant so, by Lemma 3.6.4, there is a tree Morley sequence over M, $(b_i, c_i)_{i\in\mathbb{Z}}$ so that

(a) If $i \leq j$, then $b_i c_j \equiv_M bc$.

(b) If $i > j$, then $b_i \models q|_{Mc_j}$.

Then both $(b_{2i}, c_{2i+1})_{i\in\mathbb{Z}}$ and $(b_{2i}, c_{2i-1})_{i\in\mathbb{Z}}$ are tree Morley sequences over M by Lemma 3.5.9. By (a), we know $p_0(x; b_0) \cup p_1(x; c_1)$ Kim-forks over M so

$$\bigcup_{i\in\mathbb{Z}} p_0(x; b_{2i}) \cup p_1(x; c_{2i+1})$$

is inconsistent. However, because $b_0 \mathop{\smile\!\!\!\!\!\mid}^u_M c_{-1}$ by (2), Proposition 3.6.1 gives that $p_0(x; b_0) \cup p_1(x; c_{-1})$ does not Kim-fork over M. Therefore

$$\bigcup_{i\in\mathbb{Z}} p_0(x; b_{2i}) \cup p_1(x; c_{2i-1})$$

is consistent. And this is a contradiction, as these two partial types are the same. This completes the proof. □

Corollary 3.6.6. *Suppose T is NSOP$_1$, $M \models T$, $b \equiv_M b'$ and $b \mathop{\smile\!\!\!\!\!\mid}^K_M b'$. Then there is a tree Morley sequence $(b_i)_{i<\omega}$ over M, with $b_0 = b$ and $b_1 = b'$.*

Proof. Let $p(x; b) = \text{tp}(b'/Mb)$. By induction on ordinals $\alpha \geq 1$, we will build trees $(b^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$ spread out and s-indiscernible over M so that

1. $\nu \lhd \eta$ then $b^\alpha_\nu b^\alpha_\eta \equiv_M b'b$.
2. If $1 \leq \beta < \alpha$, then $b^\alpha_{\iota_{\beta\alpha}(\eta)} = b^\beta_\eta$.

To start, let $\bar{b} = (b_i)_{i<\omega}$ be an M-invariant Morley sequence—as $b \mathop{\smile\!\!\!\!\!\mid}^K_M b'$, we may assume this sequence is Mb'-indiscernible. Define $(b^1_\eta)_{\eta\in\mathcal{T}_1}$ by $b^1_\emptyset = b'$ and $b^1_{\langle i\rangle} = b_i$. Then $(b^1_\eta)_{\eta\in\mathcal{T}_1}$ is spread out and s-indiscernible over M and clearly satisfies (1).

Now suppose given $(b^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$. Let $\langle (b^\alpha_{\eta,i})_{\eta\in\mathcal{T}_\alpha} : i < \omega\rangle$ be an M-invariant Morley sequence with $(b^\alpha_{\eta,0})_{\eta\in\mathcal{T}_\alpha} = (b^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$. Choose $b'' \mathop{\smile\!\!\!\!\!\mid}^K_M (b^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$ with

$$b'' \models \bigcup_{\eta\in\mathcal{T}_\alpha} p(x; b^\alpha_\eta),$$

(this is possible by Remark 3.6.3). By the chain condition, we may assume the sequence $\langle (b^{\alpha}_{\eta,i})_{\eta \in \mathcal{T}_\alpha} : i < \omega \rangle$ is Mb''-indiscernible and that $b'' \mathop{\smile\!\!\!\!\!\!\mid}^K_M (b^{\alpha}_{\eta,i})_{\eta \in \mathcal{T}_\alpha, i<\omega}$. Define a tree $(c_\eta)_{\eta \in \mathcal{T}_{\alpha+1}}$ by $c_\emptyset = b''$ and $c_{\langle i \rangle \frown \eta} = b^{\alpha}_{\eta,i}$. Then let $(b^{\alpha+1}_\eta)_{\eta \in \mathcal{T}_{\alpha+1}}$ be a tree which is s-indiscernible over M and locally based on $(c_\eta)_{\eta \in \mathcal{T}_{\alpha+1}}$. By an automorphism, we may assume that $b^{\alpha+1}_{\iota_{\alpha\alpha+1}(\eta)} = b^{\alpha}_\eta$ for all $\eta \in \mathcal{T}_\alpha$. This satisfies the requirements.

Finally, if δ is a limit and we are given $(b^{\alpha}_\eta)_{\eta \in \mathcal{T}_\alpha}$ for all $\alpha < \delta$, define $(b^{\delta}_\eta)_{\eta \in \mathcal{T}_\delta}$ as follows: if $\eta \in \mathcal{T}_\delta$, choose any $\alpha < \delta$ and $\nu \in \mathcal{T}_\alpha$ so that $\eta = \iota_{\alpha\delta}(\nu)$. Then define $b^{\delta}_\eta = b^{\alpha}_\nu$. By the coherence condition, this is well-defined and clearly satisfies the requirements.

To conclude, let κ be big enough for Erdős-Rado and consider $(b^{\kappa}_\eta)_{\eta \in \mathcal{T}_\kappa}$ given by the above construction. Apply Lemma 3.5.10 to find $(c_\eta)_{\eta \in \mathcal{T}_\omega}$, a Morley tree over M, based on this tree. By an automorphism, we may assume $c_{\zeta_0} = b$ and $c_{\zeta_1} = b'$. The sequence $(c_{\zeta_i})_{i<\omega}$ is the desired tree Morley sequence. □

3.7 Forking and witnesses

Basic properties of forking

Definition 3.7.1.
1. The formula $\varphi(x;b)$ *divides* over A if there is an A-indiscernible sequence $\langle b_i : i < \omega \rangle$ with $b_0 = b$ so that $\{\varphi(x;b_i) : i < \omega\}$ is inconsistent. A type $p(x)$ divides over A if it implies some formula that divides over A. Write $a \mathop{\smile\!\!\!\!\!\!\mid}^d_A B$ to mean that $\mathrm{tp}(a/AB)$ does not divide over A.

2. The formula $\varphi(x;b)$ *forks* over A if $\varphi(x;b)$ implies a finite disjunction $\bigvee_i \psi_i(x;c_i)$ where each $\psi_i(x;c_i)$ divides over A. A type $p(x)$ forks over A if it implies a formula which forks over A. We write $a \mathop{\smile\!\!\!\!\!\!\mid}^f_A B$ to mean that $\mathrm{tp}(a/AB)$ does not fork over A.

The following facts about forking and dividing are easy and well-known – see, e.g., [GIL02] [Adl05].

Fact 3.7.2. The following are true with respect to an arbitrary theory:

1. $a \mathop{\smile\!\!\!\!\!\!\mid}^d_A b$ if and only if, given any A-indiscernible sequence $I = \langle b_i : i < \omega \rangle$ with $b = b_0$, there is $a' \equiv_{Ab} a$ so that I is Aa'-indiscernible.

2. $\mathop{\smile\!\!\!\!\!\!\mid}^f$ is an invariant ternary relation on small subsets satisfying:

   a) (Extension) If $a \mathop{\smile\!\!\!\!\!\!\mid}^f_A b$, then, for all c, there is $a' \equiv_{Ab} a$ so that $a' \mathop{\smile\!\!\!\!\!\!\mid}^f_A bc$.

   b) (Base Monotonicity) If $a \mathop{\smile\!\!\!\!\!\!\mid}^f_A bc$ then $a \mathop{\smile\!\!\!\!\!\!\mid}^f_{Ab} c$.

   c) (Left Transitivity) If $a \mathop{\smile\!\!\!\!\!\!\mid}^f_{Ab} c$ and $b \mathop{\smile\!\!\!\!\!\!\mid}^f_A c$ then $ab \mathop{\smile\!\!\!\!\!\!\mid}^f_A c$.

3. For any model M,
$$a \mathop{\smile\!\!\!\!\!\!\mid}^i_M b \implies a \mathop{\smile\!\!\!\!\!\!\mid}^f_M b \implies a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b.$$

Remark 3.7.3. $\mathop{\smile\!\!\!\!\!\!|}^d$ may fail to satisfy (2)(a) in an arbitrary theory, but always satisfies (2)(b) and (2)(c).

As a warm-up to the theorem in the next subsection, we note that these properties easily give a weak form of transitivity for $\mathop{\smile\!\!\!\!\!\!|}^K$:

Lemma 3.7.4. *Suppose $a \mathop{\smile\!\!\!\!\!\!|}^d_M bc$ and $b \mathop{\smile\!\!\!\!\!\!|}^K_M c$. Then $ab \mathop{\smile\!\!\!\!\!\!|}^K_M c$.*

Proof. Assume $a \mathop{\smile\!\!\!\!\!\!|}^d_M bc$ and $b \mathop{\smile\!\!\!\!\!\!|}^K_M c$. As $b \mathop{\smile\!\!\!\!\!\!|}^K_M c$, for any M-invariant Morley sequence $I = (c_i)_{i<\omega}$ with $c_0 = c$, there is $I' = (c_i')_{i<\omega}$ with $I' \equiv_{Mc_0} I$ which is, moreover, Mb-indiscernible. By base monotonicity of $\mathop{\smile\!\!\!\!\!\!|}^d$, $a \mathop{\smile\!\!\!\!\!\!|}^d_{Mb} c$ so there is an Mab-indiscernible sequence $I'' = (c_i'')_{i<\omega}$ with $I'' \equiv_{Mbc} I'$. Thus I'' is an M-invariant Morley sequence with $c_0'' = c$ which is Mab-indiscernible. By an automorphism, we obtain $a'b' \equiv_{Mc} ab$ so that I is $Ma'b'$-indiscernible. As I was an arbitrary M-invariant Morley sequence over M, it follows that $ab \mathop{\smile\!\!\!\!\!\!|}^K_M c$. □

Morley Sequences

Definition 3.7.5. Suppose $M \models T$. An $\mathop{\smile\!\!\!\!\!\!|}^K$-Morley sequence over M is an M-indiscernible sequence $\langle b_i : i < \omega \rangle$ satisfying $b_i \mathop{\smile\!\!\!\!\!\!|}^K_M b_{<i}$. Likewise, an $\mathop{\smile\!\!\!\!\!\!|}^f$-Morley sequence over M is an M-indiscernible sequence $\langle b_i : i < \omega \rangle$ satisfying $b_i \mathop{\smile\!\!\!\!\!\!|}^f_M b_{<i}$.

Lemma 3.7.6. *Suppose the complete theory T is NSOP$_1$, $M \models T$, and $\varphi(x; b)$ does not Kim-divide over M. Then for any $\mathop{\smile\!\!\!\!\!\!|}^K$-Morley sequence $\langle b_i : i < \omega \rangle$ over M with $b_0 = b$, $\{\varphi(x; b_i) : i < \omega\}$ is non-Kim-forking over M. In particular, this set of formulas is consistent.*

Proof. By induction on n, we will show that $\{\varphi(x; b_i) : i \leq n\}$ is non-Kim-forking over M. The case of $n = 0$ follows by hypothesis. Now suppose $\{\varphi(x; b_i) : i \leq n\}$ is non-Kim-forking over M. Fix $\sigma \in \mathrm{Aut}(\mathbb{M}/M)$ with $\sigma(b_0) = b_{n+1}$. Let $a \models \{\varphi(x; b_i) : i \leq n\}$ with $a \mathop{\smile\!\!\!\!\!\!|}^K_M b_{\leq n}$. Then $\sigma(a) \equiv_M a$ and $\models \varphi(\sigma(a); b_{n+1})$. We know $b_{n+1} \mathop{\smile\!\!\!\!\!\!|}^K_M b_{\leq n}$ so by the independence theorem, there is a' with $a' \equiv_{Mb_{\leq n}} a$ and $a' \equiv_{Mb_{n+1}} \sigma(a)$ so that $a' \mathop{\smile\!\!\!\!\!\!|}^K_M b_{\leq n+1}$. As $a' \models \{\varphi(x; b_i) : i \leq n+1\}$, this completes the induction. The lemma, then, follows by compactness. □

Theorem 3.7.7. *Suppose the complete theory T is NSOP$_1$ and $M \models T$. The following are equivalent:*

1. *$\varphi(x; b)$ Kim-divides over M.*
2. *For some $\mathop{\smile\!\!\!\!\!\!|}^f$-Morley sequence $(b_i)_{i<\omega}$ over M with $b_0 = b$, $\{\varphi(x; b_i) : i < \omega\}$ is inconsistent.*
3. *For every $\mathop{\smile\!\!\!\!\!\!|}^f$-Morley sequence $(b_i)_{i<\omega}$ over M with $b_0 = b$, $\{\varphi(x; b_i) : i < \omega\}$ is inconsistent.*

Proof. (3) $\implies$ (2) is immediate, as a Morley sequence in a global M-invariant type is, in particular, an $\mathop{\smile\!\!\!\!\!\!\mid}^f$-Morley sequence and such sequences always exist.

(2) $\implies$ (1) follows from Lemma 3.7.6, as an $\mathop{\smile\!\!\!\!\!\!\mid}^f$-Morley sequence is an $\mathop{\smile\!\!\!\!\!\!\mid}^K$-Morley sequence.

Now we show (1) $\implies$ (3). Suppose not—assume that $\varphi(x;b)$ is a formula which Kim-divides over M, but there is some $\mathop{\smile\!\!\!\!\!\!\mid}^f$-Morley sequence over M with $b_0 = b$ so that $\{\varphi(x;b_i) : i < \omega\}$ is consistent. By induction on n, we will construct a sequence $(b'_i)_{i\leq n}$ and an elementary chain $(N_i)_{i\leq n}$ so that

1. For all $n < \omega$, $b_0 \dots b_n \equiv_M b'_0 \dots b'_n$.
2. For all $n < \omega$, $M \prec N_n \prec N_{n+1} \prec \mathbb{M}$.
3. For all $n < \omega$, $b'_n \mathop{\smile\!\!\!\!\!\!\mid}^f_M N_n$.
4. For all $n < \omega$, $b'_n \in N_{n+1}$.

For the $n = 0$ case, set $b'_0 = b_0$ and $N_0 = M$. Now suppose we are given $(N_i)_{i\leq n}$ and $(b'_i)_{i\leq n}$. Let N_{n+1} be an arbitrary (small) elementary extension of N_n which contains b'_n. By invariance and extension of $\mathop{\smile\!\!\!\!\!\!\mid}^f$, we may choose some b'_{n+1} so that $b'_0 \dots b'_{n+1} \equiv_M b_0 \dots b_{n+1}$ and $b'_{n+1} \mathop{\smile\!\!\!\!\!\!\mid}^f_M N_{n+1}$. This completes the recursion.

Set $N = \bigcup_{i<\omega} N_i$.

Claim 1: For all $n < \omega$, $(b'_i)_{i\geq n} \mathop{\smile\!\!\!\!\!\!\mid}^f_M N_n$.

Proof of claim: Fix n. We will argue by induction on k that $b'_n \dots b'_{n+k} \mathop{\smile\!\!\!\!\!\!\mid}^f_M N_n$. For $k = 0$, this is by construction. Assume it has been proven for k. Note that $b'_{n+k+1} \mathop{\smile\!\!\!\!\!\!\mid}^f_M N_{n+k+1}$. Now N_n and $(b'_i)_{i\leq n+k}$ are contained in N_{n+k+1} so, in particular, we have $b'_{n+k+1} \mathop{\smile\!\!\!\!\!\!\mid}^f_M N_n b'_0 \dots b'_{n+k}$. By base monotonicity, we have

$$b'_{n+k+1} \mathop{\smile\!\!\!\!\!\!\mid}^f_{M b'_0 \dots b'_{n+k}} N_n.$$

This, together with the induction hypothesis, implies

$$b'_0 \dots b'_{n+k+1} \mathop{\smile\!\!\!\!\!\!\mid}^f_M N_n$$

by left-transitivity. The claim follows by finite character. □

Let $\mathcal{D}$ be any non-principal ultrafilter on $\{b'_i : i < \omega\}$ and $(c_i)_{i<\omega}$ be a sequence chosen so that $c_i \models \mathrm{Av}(\mathcal{D}, Nc_{<i})$, i.e. a Morley sequence over N in the global $(b'_i)_{i<\omega}$-invariant type $\mathrm{Av}(\mathcal{D}, \mathbb{M})$.

Claim 2: $(c_i)_{i<\omega} \mathop{\smile\!\!\!\!\!\!\mid}^f_M N$.

Proof of claim: Suppose not. Then by finite character, there is l so that $(c_i)_{i<l} \not\mathop{\smile\!\!\!\!\!\!\mid}^f_M N$ so we choose some $\varphi(x_0, \dots, x_{l-1}; d) \in \mathrm{tp}(c_0, \dots, c_{l-1}/N)$ which forks over M. Choose n

so that $d \in N_n$. By definition of average type, we may find $i_0 > \ldots > i_{l-1} > n$ so that $\mathbb{M} \models \varphi(b'_{i_0}, \ldots, b'_{i_{l-1}}; d)$. Then $(b'_i)_{i \geq n} \not\mathrel{\unicode{x2AEB}}^f_M N_n$, contradicting Claim 1. □

Let $q \supseteq \mathrm{tp}((c_i)_{i<\omega}/M)$ be a global M-invariant and indiscernible type, as in Definition 3.3.7. Let $\langle (c_{k,i})_{i<\omega} : k < \omega \rangle$ be a Morley sequence over M in q with $c_{0,i} = c_i$ for all $i < \omega$. By Lemma 3.3.10, $(\overline{c}_k)_{k<\omega}$ is a mutually-indiscernible array over M. By Claim 2, we know $\overline{c}_0 \mathrel{\unicode{x2AEB}}^f_M N$ hence $\overline{c}_0 \mathrel{\unicode{x2AEB}}^K_M N$, so we may assume the sequence $(\overline{c}_k)_{k<\omega}$ is N-indiscernible by symmetry. We know that $\{\varphi(x; b_i) : i < \omega\}$ is consistent so $\{\varphi(x; b'_i) : i < \omega\}$ is consistent, and therefore $\{\varphi(x; c_{0,i}) : i < \omega\}$ is consistent. The sequence $(c_{0,i})_{i<\omega}$ is also an N-invariant Morley sequence so $\varphi(x; c_{0,0})$ does not Kim-divide over N. But as $c_{0,0} \equiv_M b$, $(c_{i,0})_{i<\omega}$ is an M-invariant Morley sequence over M, and $\varphi(x; b)$ Kim-divides over M, we know that $\{\varphi(x; c_{i,0}) : i < \omega\}$ is inconsistent.

Let $(\overline{d}_i)_{i<\omega}$ be a mutually indiscernible array over N, locally based on $(\overline{c}_i)_{i<\omega}$ (exists by [Che14, Lemma 1.2]), with $(\overline{d}_i)_{i<\omega}$ an N-indiscernible sequence. By Lemma 3.3.10, we have $(d_{i,0})_{i<\omega} \equiv_M (c_{i,0})_{i<\omega}$. Also, because $(\overline{c}_i)_{i<\omega}$ was taken to be N-indiscernible and $\overline{c}_0$ was an N-invariant Morley sequence, we know each $\overline{c}_i$ is an N-invariant Morley sequence, and therefore each $\overline{d}_i$ is an N-invariant Morley sequence. By choice of the array, $\{\varphi(x; d_{i,j}) : j < \omega\}$ is consistent for all i, so $\varphi(x; d_{i,0})$ does not Kim-divide over N. Also, we have $\{\varphi(x; d_{i,0}) : i < \omega\}$ is inconsistent. Thus, to derive a contradiction, it suffices by Lemma 3.7.6 to establish the following:

Claim 3: $(d_{i,0})_{i<\omega}$ is an $\mathrel{\unicode{x2AEB}}^K$-Morley sequence over N.

Proof of claim: As the $(d_{i,j})_{i,j<\omega}$ forms a mutually indiscernible array over N, we know that for each $i < \omega$, $\overline{d}_i$ is an $N\overline{d}_{<i}$-indiscernible sequence. But it is also an N-invariant Morley sequence so $\overline{d}_{<i} \mathrel{\unicode{x2AEB}}^K_N d_{i,0}$. By symmetry, this yields in particular that $d_{i,0} \mathrel{\unicode{x2AEB}}^K_N d_{0,0} \ldots d_{i-1,0}$. This proves the claim and completes the proof. □

Witnesses

Definition 3.7.8. Suppose M is a model and $(a_i)_{i<\omega}$ is an M-indiscernible sequence.

1. Say $(a_i)_{i<\omega}$ is a *witness* for Kim-dividing over M if, whenever $\varphi(x; a_0)$ Kim-divides over M, $\{\varphi(x; a_i) : i < \omega\}$ is inconsistent.

2. Say $(a_i)_{i<\omega}$ is a *strong witness* to Kim-dividing over M if, for all n, the sequence $\langle (a_{n \cdot i}, a_{n \cdot i+1}, \ldots, a_{n \cdot i+n-1}) : i < \omega \rangle$ is a witness to Kim-dividing over M.

Corollary 3.5.14 and Lemma 3.5.9 show that tree Morley sequences are strong witnesses for Kim-dividing. The following proposition shows the converse, giving a characterization of strong witnesses as exactly the tree Morley sequences.

Proposition 3.7.9. *Suppose T is NSOP$_1$ and $M \models T$. Then $(a_i)_{i<\omega}$ is a strong witness for Kim-dividing over M if and only if $(a_i)_{i<\omega}$ is a tree Morley sequence over M.*

Proof. If $(a_i)_{i<\omega}$ is a tree Morley sequence, then $(a_{n\cdot i}, a_{n\cdot i+1}, \ldots, a_{n\cdot i+(n-1)})_{i<\omega}$ is also a tree Morley sequence over M by Lemma 3.5.9. It follows that $(a_i)_{i<\omega}$ is a strong witness to Kim-dividing by Corollary 3.5.14.

For the other direction, suppose $(a_i)_{i<\omega}$ is a strong witness to Kim-dividing over M. Given an arbitrary cardinal κ, we may, by compactness, stretch the sequence to $(a_i)_{i<\kappa}$ which is still a strong witness to Kim-dividing over M. By recursion on $\alpha < \kappa$, we will construct trees $(a^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$ so that

1. For all $i < \alpha$, $a^\alpha_{\zeta_i} = a_i$ and also $a^\alpha_\emptyset = a_\alpha$ for α successor.
2. $(a^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$ is spread out over M and s-indiscernible over $M(a_i)_{i>\alpha}$.
3. If $\alpha < \beta$, then $a^\alpha_\eta = a^\beta_{\iota_{\alpha\beta}(\eta)}$ for all $\eta \in \mathcal{T}_\alpha$.

For the case $\alpha = 0$, put $a^0_\emptyset = a_0$. This satisfies the demands. Suppose $(a^\beta_\eta)_{\eta\in\mathcal{T}_\beta}$ has been defined for all $\beta \le \alpha$. By Ramsey, compactness, and an automorphism, we may assume $(a_i)_{i>\alpha}$ is $M(a^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$-indiscernible. As $I_{>\alpha} = (a_i)_{i>\alpha}$ is also a strong witness to Kim-dividing over M, we have

$$(a^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha} \mathop{\downarrow}^K_M I_{>\alpha}.$$

Let $J = \langle (a^\alpha_{\eta,i})_{\eta\in\mathcal{T}_\alpha} : i < \omega\rangle$ be a Morley sequence in an M-invariant type with $a^\alpha_{\eta,0} = a^\alpha_\eta$ for all $\eta \in \mathcal{T}_\alpha$. By symmetry, $I_{>\alpha} \mathop{\downarrow}^K_M (a^\alpha_\eta)_{\eta\in\mathcal{T}_\alpha}$ so we may assume J is $MI_{>\alpha}$-indiscernible. Define the tree $(a^{\alpha+1}_\eta)_{\eta\in\mathcal{T}_{\alpha+1}}$ by $a^{\alpha+1}_\emptyset = a_{\alpha+1}$ and $a^{\alpha+1}_{\langle i\rangle\frown\eta} = a^\alpha_{\eta,i}$ for all $\eta \in \mathcal{T}_\alpha$ and $i < \omega$. Note in particular, this definition gives $a^{\alpha+1}_{\iota_{\alpha\alpha+1}(\eta)} = a^{\alpha+1}_{0\frown\eta} = a^\alpha_\eta$ for all $\eta \in \mathcal{T}_\alpha$. The tree we just constructed is clearly spread out. By an automorphism, we may further assume $(a^{\alpha+1}_\eta)_{\eta\in\mathcal{T}_{\alpha+1}}$ is s-indiscernible over $MI_{>\alpha+1}$. This completes the successor step.

Now suppose given $(a^\beta_\eta)_{\eta\in\mathcal{T}_\beta}$ for all $\beta < \delta$, where δ is a limit. Define $(a^\delta_\eta)_{\eta\in\mathcal{T}_\delta}$ by setting $a^\delta_{\iota_{\alpha\delta}(\eta)} = a^\alpha_\eta$ for all $\alpha < \delta$ and $\eta \in \mathcal{T}_\alpha$. Condition (3) guarantees that this is well-defined.

Taking κ to be sufficiently large, we may extract a Morley tree from the tree we just constructed by Lemma 3.5.10 – in particular, we may obtain a Morley tree $(b_\eta)_{\eta\in\mathcal{T}_\omega}$ so that $(b_{\zeta_i})_{i<\omega} \equiv_M (a_i)_{i<\omega}$. This shows that $(a_i)_{i<\omega}$ is a tree Morley sequence over M. □

Corollary 3.7.10. *Suppose T is NSOP$_1$ and $M \models T$. An $\mathop{\downarrow}^f$-Morley sequence over M is a tree Morley sequence.*

Proof. Suppose $(a_i)_{i<\omega}$ is an $\mathop{\downarrow}^f$-Morley sequence over M. Arguing as in Claim 1 of the proof of Theorem 3.7.7, for all $n < \omega$, $a_{>n} \mathop{\downarrow}^f_M a_{\le n}$. Therefore, $\langle (a_{n\cdot i}, a_{n\cdot i+1}, \ldots, a_{n\cdot i+n-1}) : i < \omega\rangle$ is an $\mathop{\downarrow}^f$-Morley sequence over M, hence a witness to Kim-dividing over M by Theorem 3.7.7. This shows $(a_i)_{i<\omega}$ is a strong witness to Kim-dividing over M. By Proposition 3.7.9, $(a_i)_{i<\omega}$ is a tree Morley sequence over M. □

In any theory, if $(a_i)_{i<\omega}$ is an $\mathop{\smile\!\!\!\!\!\!|}^f$-Morley sequence over A, then, as the proof of Corollary 3.7.10 shows, that $a_{>n} \mathop{\smile\!\!\!\!\!\!|}^f_A a_{\leq n}$ for all $n < \omega$. As base monotonicity and left-transitivity do not necessarily hold for $\mathop{\smile\!\!\!\!\!\!|}^K$, we give a Morley sequence with this stronger behavior a name:

Definition 3.7.11. Say the M-indiscernible sequence $(a_i)_{i<\omega}$ is a *total* $\mathop{\smile\!\!\!\!\!\!|}^K$*-Morley sequence* if $a_{>n} \mathop{\smile\!\!\!\!\!\!|}^K_M a_{\leq n}$ for all $n < \omega$.

3.8 Characterizing NSOP$_1$ and simple theories

The Main Theorem

Before continuing with the rest of the paper, we pause to take stock of what has been shown:

Theorem 3.8.1. *The following are equivalent for the complete theory T:*

1. *T is NSOP$_1$*

2. *Ultrafilter independence of higher formulas: for every model $M \models T$, and ultrafilters $\mathcal{D}$ and $\mathcal{E}$ on M with $Av(\mathcal{D}, M) = Av(\mathcal{E}, M)$, $(\varphi, M, \mathcal{D})$ is higher if and only if $(\varphi, M, \mathcal{E})$ is higher*

3. *Kim's lemma for Kim-dividing: For every model $M \models T$ and $\varphi(x; b)$, if $\varphi(x; y)$ q-divides for some global M-invariant $q \supseteq tp(b/M)$, then $\varphi(x; y)$ q-divides for* every *global M-invariant $q \supseteq tp(b/M)$.*

4. *Local character: for some infinite cardinal κ, there cannot be a sequence $\langle N_i, \varphi_i(x, y_i), c_i : i < \kappa \rangle$ such that $\langle N_i : i < \kappa \rangle$ is an increasing continuous sequence of models of T, $\varphi_i(x, y_i)$ is a formula over N_i, $c_i \in N_{i+1}$, such that $\varphi_i(x, c_i)$ Kim-forks over N_i and $\{\varphi(x, c_i) : i < \kappa\}$ is consistent.*

5. *Symmetry over models: for every $M \models T$, then $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$ if and only if $b \mathop{\smile\!\!\!\!\!\!|}^K_M a$.*

6. *Independence theorem over models: if $M \models T$, $a \equiv_M a'$, $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$, $a' \mathop{\smile\!\!\!\!\!\!|}^K_M c$, and $b \mathop{\smile\!\!\!\!\!\!|}^K_M c$, then there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\!\!\!\!\!\!|}^K_M bc$.*

Proof. (1) $\iff$ (2) $\iff$ (3) is Theorem 3.3.16.
(1) $\iff$ (4) is Corollary 3.4.6.
(1) $\iff$ (5) is Theorem 3.5.16.
(1) $\iff$ (6) is Theorem 3.6.5. $\square$

Simplicity within the class of NSOP$_1$ theories

Definition 3.8.2. [Che14, Section 6] Suppose $p(x)$ is a partial type over the set A.

1. We say p is a *simple type* if there is no $\varphi(x;y)$, $(a_\eta)_{\eta\in\omega^{<\omega}}$ and $k<\omega$ so that $\{\varphi(x;a_{\eta\frown\langle i\rangle}) : i<\omega\}$ is k-inconsistent for all $\eta\in\omega^{<\omega}$ and $p(x)\cup\{\varphi(x;a_{\eta|i}) : i<\omega\}$ is consistent for all $\eta\in\omega^\omega$. Equivalently, $p(x)$ is simple if, whenever $B\supseteq A$, $q\in S(B)$, and $p\subseteq q$, then q does not divide over AB' for some $B'\subseteq B$, $|B'|\leq|T|$ (for the definition of dividing, see Definition 3.7.1 above).
2. We say $p(x)$ is a *co-simple type* if there is no formula $\varphi(x;y)\in L(A)$ for which there exists $(a_\eta)_{\eta\in\omega^{<\omega}}$ and $k<\omega$ so that $\{\varphi(x;a_{\eta\frown\langle i\rangle}) : i<\omega\}$ is k-inconsistent for all $\eta\in\omega^{<\omega}$ and $\{\varphi(x;a_{\eta|i}) : i<\omega\}$ is consistent for all $\eta\in\omega^\omega$ and moreover $a_\eta\models p$ for all $\eta\in\omega^{<\omega}$.

Proposition 3.8.3. *Assume T is NSOP$_1$ and let $\pi(x)$ be a partial type over A.*

1. *Assume that for any $\varphi(x;a)$ and any model $M\supseteq A$, $\pi(x)\cup\{\varphi(x;a)\}$ divides over M if and only if $\pi\cup\{\varphi(x;a)\}$ Kim-divides over M. Then $\pi(x)$ is a simple type.*
2. *Assume that if $M\supseteq A$, then for any a and for any $b\models\pi(x)$, $a\mathop{\smile\!\!\!\!\!\!\mid}^f_M b$ if and only if $a\mathop{\smile\!\!\!\!\!\!\mid}^K_M b$. Then π is a co-simple type.*

Proof. Fix a Skolemization T^{Sk} of T. Throughout the proof, indiscernibility will be with respect to the language L^{Sk} of the Skolemization. (1) Suppose π is not simple. Then by compactness, there is a formula $\varphi(x;y)$ over A and a tree $(a_\eta)_{\eta\in\omega^{<\omega+1}}$ s-indiscernible over A so that for some $k<\omega$

- For all $\eta\in\omega^{\omega+1}$, $\pi(x)\cup\{\varphi(x;a_{\eta|\alpha}) : \alpha<\omega+1\}$ is consistent
- For all $\eta\in\omega^{<\omega+1}$, $\{\varphi(x;a_{\eta\frown\alpha}) : \alpha<\omega\}$ is k-inconsistent.

Moreover we may assume $(a_{0^\alpha} : \alpha<\omega+1)$ is an A-indiscernible sequence. Let $b\models\pi(x)\cup\{\varphi(x;a_{0^\alpha}) : \alpha<\omega+1\}$. By Ramsey, compactness, and automorphism, we may assume $(a_{0^\alpha} : \alpha<\omega+1)$ is Ab-indiscernible. Let $C=\{a_{0^\alpha} : \alpha<\omega\}$. Then s-indiscernibility implies $(a_{0^\omega\frown\beta} : \beta<\omega)$ is indiscernible over $A\cup C$ and $\{\varphi(x;a_{0^\omega\frown\beta}) : \beta<\omega\}$ is k-inconsistent by our assumption. As $b\models\varphi(x;a_{0^\omega\frown\langle 0\rangle})$, we have $b\not\mathop{\smile\!\!\!\!\!\!\mid}^d_{AC} a_{0^\omega\frown\langle 0\rangle}$. But by indiscernibility, $a_{0^\omega\frown\langle 0\rangle}\mathop{\smile\!\!\!\!\!\!\mid}^u_{AC} b$ so in particular $a_{0^\omega\frown\langle 0\rangle}\mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ and $b\mathop{\smile\!\!\!\!\!\!\mid}^K_M a_{0^\omega\frown\langle 0\rangle}$, where $M=Sk(AC)$, by symmetry.

(2) We argue similarly. Suppose $(a_\eta)_{\eta\in\omega^{<\omega+1}}$ is a collection of realizations of π, forming a tree s-indiscernible over A, with respect to which $\varphi(x;y)$ witnesses that π is not co-simple. Let $a\models\{\varphi(x;b_{0^\alpha}) : \alpha<\omega+1\}$. By Ramsey, compactness, and automorphism, we may assume $(b_{0^\alpha} : \alpha<\omega+1)$ is a Ba-indiscernible sequence. Setting $M=\mathrm{Sk}(A(b_{0^\alpha})_{\alpha<\omega})$, we have $a\not\mathop{\smile\!\!\!\!\!\!\mid}^d_M b_{0^\omega\frown\langle 0\rangle}$ but $b_{0^\omega\frown\langle 0\rangle}\mathop{\smile\!\!\!\!\!\!\mid}^u_M a$ so $a\mathop{\smile\!\!\!\!\!\!\mid}^K_M b_{0^\omega\frown\langle 0\rangle}$. $\square$

In a similar vein, we have:

Proposition 3.8.4. *The complete theory T is simple if and only if T is NSOP$_1$ and $\mathop{\smile\!\!\!\!\!\mid}^f = \mathop{\smile\!\!\!\!\!\mid}^K$ over models.*

Proof. If T is simple, then $\mathop{\smile\!\!\!\!\!\mid}^f = \mathop{\smile\!\!\!\!\!\mid}^K$ over models by Kim's lemma for simple theories [Kim98, Proposition 2.1], as a Morley sequence in a global invariant type is, in particular, a Morley sequence in the sense of non-forking. On the other hand, by [Kim01, Theorem 2.4] forking is symmetric if and only if T is simple and, by [Che14, Lemma 6.16], we even have that if forking is symmetric over models then T is simple. If T is NSOP$_1$, then $\mathop{\smile\!\!\!\!\!\mid}^K$ is symmetric so $\mathop{\smile\!\!\!\!\!\mid}^K = \mathop{\smile\!\!\!\!\!\mid}^f$ implies T is simple. $\square$

We also can give an interesting new proof of the following well-known fact:

Corollary 3.8.5. *The complete theory T is simple if and only if T is NSOP$_1$ and NTP$_2$.*

Proof. In an NTP$_2$ theory, if $\varphi(x; b)$ divides over a model M, there is a Morley sequence sequence over M in some global M-finitely satisfiable type witnessing this [CK12, Lemma 3.14]. So $\mathop{\smile\!\!\!\!\!\mid}^d = \mathop{\smile\!\!\!\!\!\mid}^K$, which implies T is simple. $\square$

Definition 3.8.6. [YC14, Definition 2.5] We say $(a_i)_{i\in\kappa}$ is a *universal Morley sequence* in $p \in S(A)$ if

- $(a_i)_{i\in\kappa}$ is indiscernible with $a_i \models p$
- If $\varphi(x; y) \in L(A)$ and $\varphi(x; a_0)$ divides over A then $\{\varphi(x; a_i) : i \in \kappa\}$ is inconsistent.

Proposition 3.8.7. *Suppose T is NSOP$_1$. Then T is simple if and only if, for any $M \models T$ and $p(x) \in S(M)$, there is a universal Morley sequence in p.*

Proof. If T is simple, then in any type $p(y) \in S(M)$, there is a $\mathop{\smile\!\!\!\!\!\mid}^f$-Morley sequence in $p(y)$. By Kim's lemma for simple theories [Kim98, Proposition 2.1], this is a universal Morley sequence in p.

If T is not simple, then there is some formula $\varphi(x; b) \in L(Mb)$ which divides over M but does not Kim-divide over M, by Proposition 3.8.4. Suppose there is a universal Morley sequence in $\text{tp}(b/M)$—by compactness we can take it to be $(b_i)_{i\in\mathbb{Z}}$ indexed by $\mathbb{Z}$. Then given $i \in \mathbb{Z}$, we have $b_{<i}$ is Mb_i-indiscernible so $b_{<i} \mathop{\smile\!\!\!\!\!\mid}^d_M b_i$ so $b_i \mathop{\smile\!\!\!\!\!\mid}^K_M b_{<i}$ by symmetry. So $(b_i)_{i\in\mathbb{Z}}$ is an $\mathop{\smile\!\!\!\!\!\mid}^K$-Morley sequence. By Lemma 3.7.6, $\{\varphi(x; b_i) : i \in \mathbb{Z}\}$ is consistent. But $\varphi(x; b)$ divides over M and $(b_i)_{i\in\mathbb{Z}}$ is a universal Morley sequence so $\{\varphi(x; b_i) : i \in \mathbb{Z}\}$ is inconsistent. This is a contradiction. $\square$

If $a \mathop{\smile\!\!\!\!\!\mid}^K_M bb'$, it does not always make sense to ask if $a \mathop{\smile\!\!\!\!\!\mid}^K_{Mb} b'$, since it is not always the case that $\text{tp}(b'/Mb)$ extends to a global Mb'-invariant type. This can occur, however, whenever Mb' is a model, for instance. Say $\mathop{\smile\!\!\!\!\!\mid}^K$ satisfies base monotonicity *over models* if, whenever $a \mathop{\smile\!\!\!\!\!\mid}^K_M Nb$ where $M, N \models T$, then $a \mathop{\smile\!\!\!\!\!\mid}^K_N b$.

Proposition 3.8.8. *The NSOP$_1$ theory T is simple if and only if $\mathop{\smile\!\!\!\!\!\!|}^K$ satisfies base monotonicity over models.*

Proof. It T is simple, this follows from Proposition 3.8.4, using Fact 3.7.2(2b). On the other hand, suppose $\mathop{\smile\!\!\!\!\!\!|}^K$ satisfies base monotonicity over models. We will show that $\mathop{\smile\!\!\!\!\!\!|}^K = \mathop{\smile\!\!\!\!\!\!|}^d$ over models. It follows then that T is simple, by Proposition 3.8.4. So suppose towards contradiction that $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$ but $a \not\mathop{\smile\!\!\!\!\!\!|}^d_M b$, witnessed by $\varphi(x; b) \in \text{tp}(a/Mb)$ and $I = (b_i)_{i<\omega+1}$ an M-indiscernible sequence with $b_\omega = b$ and $\{\varphi(x; b_i) : i < \omega + 1\}$ inconsistent. Fix a Skolemization T^{Sk} of T. By Ramsey and automorphism, we may assume $(b_i : i < \omega + 1)$ is L^{Sk}-indiscernible over M. As $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$, we may, by extension, assume $a \mathop{\smile\!\!\!\!\!\!|}^K_M \text{Sk}(MI)$. Let $N = Sk(MI_{<\omega})$. By base monotonicity over models, we have $a \mathop{\smile\!\!\!\!\!\!|}^K_N b$. But stretching I to $(b_i)_{i<\omega+\omega}$, we have that $(b_{\omega+i})_{i<\omega}$ is a N-invariant Morley sequence (in the reverse order) in $\text{tp}(b/N)$ and $\{\varphi(x; b_{\omega+i}) : i < \omega\}$ is inconsistent. So $a \not\mathop{\smile\!\!\!\!\!\!|}^K_N b$, a contradiction. □

3.9 Examples

A Kim-Pillay-style characterization of $\mathop{\smile\!\!\!\!\!\!|}^K$

We are interested in explicitly describing $\mathop{\smile\!\!\!\!\!\!|}^K$ in concrete examples. As in simple theories, this is most easily acheived by establishing the existence of an independence relation with certain properties and then deducing that, therefore, the relation coincides with $\mathop{\smile\!\!\!\!\!\!|}^K$. The following theorem explains how this works. The content of the theorem is essentially the same as Proposition I.5.8, where a Kim-Pillay style criterion for NSOP$_1$ theories was observed, but we point out how this gives information about Kim-independence.

Theorem 3.9.1. *Assume there is an $Aut(\mathbb{M})$-invariant ternary relation $\mathop{\smile\!\!\!\!\!\!|}$ on small subsets of the monster $\mathbb{M} \models T$ which satisfies the following properties, for an arbitrary $M \models T$ and arbitrary tuples from $\mathbb{M}$.*

1. *Strong finite character: if $a \not\mathop{\smile\!\!\!\!\!\!|}_M b$, then there is a formula $\varphi(x, b, m) \in tp(a/bM)$ such that for any $a' \models \varphi(x, b, m)$, $a' \not\mathop{\smile\!\!\!\!\!\!|}_M b$.*
2. *Existence over models: $M \models T$ implies $a \mathop{\smile\!\!\!\!\!\!|}_M M$ for any a.*
3. *Monotonicity: $aa' \mathop{\smile\!\!\!\!\!\!|}_M bb' \implies a \mathop{\smile\!\!\!\!\!\!|}_M b$.*
4. *Symmetry: $a \mathop{\smile\!\!\!\!\!\!|}_M b \iff b \mathop{\smile\!\!\!\!\!\!|}_M a$.*
5. *The independence theorem: $a \mathop{\smile\!\!\!\!\!\!|}_M b$, $a' \mathop{\smile\!\!\!\!\!\!|}_M c$, $b \mathop{\smile\!\!\!\!\!\!|}_M c$ and $a \equiv_M a'$ implies there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\!\!\!\!\!\!|}_M bc$*

Then T is NSOP$_1$ and $\mathop{\smile\!\!\!\!\!\!|}$ strengthens *$\mathop{\smile\!\!\!\!\!\!|}^K$—i.e. if $M \models T$, $a \mathop{\smile\!\!\!\!\!\!|}_M b$ then $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$. If, moreover, $\mathop{\smile\!\!\!\!\!\!|}$ satisfies*

6. *Witnessing: if $a \not\perp_M b$ witnessed by $\varphi(x; b)$ and $(b_i)_{i<\omega}$ is a Morley sequence over M in a global M-invariant type extending tp(b/M), then $\{\varphi(x; b_i) : i < \omega\}$ is inconsistent.*

then $\perp = \perp^K$.

Proof. It was shown in Proposition I.5.8 that if there is such a relation $\perp$, then T is NSOP$_1$. The proof there shows that if $\perp$ satisfies axioms (1)-(4), then $a \perp^u_M b$ implies $a \perp_M b$. Now suppose $a \perp_M b$. Let $p(x; b) = \text{tp}(a/Mb)$ and let q be a global coheir of $\text{tp}(b/M)$. By the independence theorem for $\perp$, if $(b_i)_{i<\omega}$ is a Morley sequence over M in q with $b_0 = b$, then $\bigcup_{i<\omega} p(x; b_i)$ is consistent. But then $a \perp^K_M b$. The "moreover" clause follows by definition of $\perp^K$. $\square$

Remark 3.9.2. The condition (6) can be weakened to quantifying only over global *coheirs* of $\text{tp}(b/M)$, or asserting the existence of *one* such coheir – this is sometimes slightly easier in practice.

Remark 3.9.3. Axioms (1)-(5) do not, by themselves, suffice to characterize $\perp^K$. See Remark 3.9.37 below.

Combinatorial examples

In this section, we study some combinatorial examples of NSOP$_1$ theories which are not simple. They are structures which encode a generic family of selector functions for an equivalence relation. The theories defined below provide a different presentation of a theory defined by Džamonja and Shelah in [DS04] (where it was called T^*_{feq} – though this name is now typically reserved for a different theory) and later studied by Malliaris in [Mal12] (where it was called T^s). We give a family of theories T^*_n as n ranges over positive integers, but we will only be interested in the case of $n = 1, 2$. Among non-simple NSOP$_1$ theories, the theory T^*_1 is probably the easiest to understand, and we show that already T^*_1 witnesses many of the new phenomena in our context: with respect to this theory, we give explicit examples of formulas which divide but do not Kim-divide, formulas which fork and do not divide over models, and types which contain no universal Morley sequences.

We use T^*_1 to answer a question of Chernikov from [Che14] concerning simple and co-simple types and a question of Conant from [Con14] concerning forking and dividing. A type is simple if no instance of the tree property is consistent with the type and a type is co-simple if the tree property cannot witnessed using parameters which realize the type (see Definition 3.8.2 above for the precise definition). For stability, no such distinction arises, but Chernikov was able to show that, in general, there are co-simple types which are not simple. In fact, examples can be found in the triangle-free random graph. It was asked if there can exist simple types which are not co-simple and he showed that there can be no such types in an NTP$_2$ theory. In [Con14], Conant gave a detailed analysis of forking and dividing in the Henson graphs and showed that forking does not equal dividing for formulas, though every complete type has a global non-forking extension. As the Henson graphs all have the property SOP_3, Conant asked if there could be an NSOP_3 example of this behavior. We show the answer to both questions is yes already within the class of NSOP_1 theories.

Lastly, we use T_2^* to give a counter-example to transitivity for $\mathop{\smile\!\!\!\!\!\!|}^K$. Because Kim-dividing does not behave well with respect to changing the base, the normal formulation of transitivity does not necessarily make sense. Nonetheless, there is a natural way to formulate a version which does make sense. Suppose T is NSOP_1, $M \models T$ and both $a \mathop{\smile\!\!\!\!\!\!|}^K_M bc$ and $b \mathop{\smile\!\!\!\!\!\!|}^K_M c$. Must it also be the case, then, that $ab \mathop{\smile\!\!\!\!\!\!|}^K_M c$? We show the answer is no.

For the remainder of this subsection, if A is a structure in some language and $X \subseteq A$, write $\langle X \rangle^A$ for the substructure of A generated by X. We write just $\langle X \rangle$ when A is the monster model.

For a natural number $n \geq 1$, let $L_n = \langle O, F, E, \mathrm{eval} \rangle$ where O, F are sorts, E is a binary relation symbol, and eval is an $n+1$-ary function. The theory T_n will say

- O and F are sorts—O and F disjoint and the universe is their union.
- $E \subseteq O^2$ is an equivalence relation on O.
- $\mathrm{eval} : F^n \times O \to O$ is a function so that for all $f \in F^n$, $\mathrm{eval}(f, -)$ is a function from O to O which is a selector function for E – more formally, for all $b \in O$, we have $E(\mathrm{eval}(f,b), b)$ and if $b, b' \in O$ and $E(b,b')$ then we have
$$\mathrm{eval}(f,b) = \mathrm{eval}(f,b').$$

The letter F is for 'function' and O is for 'object'—we think of a tuple $f \in F^n$ as naming the function $\mathrm{eval}(f,-)$. Let $\mathbb{K}_n$ be the class of finite models of T_n.

Recall that a Fraïssé class $\mathbb{K}$ is said to have the *strong amalgamation property (SAP)* if, whenever $A, B, C \in \mathbb{K}$, and $e : A \to B$ and $f : A \to C$ are embeddings, then there is a structure $D \in \mathbb{K}$ and embeddings $g : B \to D$, $h : C \to D$ so that $ge = hf$ and, moreover, $(\mathrm{im} g) \cap (\mathrm{im} h) = (\mathrm{im} ge)$ (and hence also $= (\mathrm{im} hf)$).

Lemma 3.9.4. *The class $\mathbb{K}_n$ is a Fraïssé class with SAP. Moreover, it is uniformly locally finite.*

Proof. HP is clear as the axioms of T_n are universal. The argument for JEP is identical to that for SAP, so we show SAP. Suppose $A, B, C \in \mathbb{K}_n$ where $A \subseteq B, C$ and $B \cap C = A$. It suffices to define a L_n-structure with domain $D = B \cup C$, extending both B and C. Interpret O^D and F^D by $O^D = O^B \cup O^C$ and $F^D = F^B \cup F^C$. Let E^D be the equivalence relation generated by $E^B \cup E^C$. It follows that if $b \in B$, $c \in C$ and $(b,c) \in E^D$, then there is some $a \in A$ so that $(a,b) \in E^B$ and $(a,c) \in E^C$ and, moreover, (O^D, E^D) extends both (O^B, E^B) and (O^C, E^C) as equivalence relations.

We are left with interpreting eval^D. Let $\{a_i : i < k_0\}$ enumerate a collection of representatives for the E^A-classes in A. Then let $\{b_i : i < k_1\}$ and $\{c_i : i < k_2\}$ enumerate representatives for the E^B- and E^C-classes of elements not represented by an element of A, respectively. Then every element of O^D is equivalent to a unique element of

$$X = \{a_i : i < k_0\} \cup \{b_i : i < k_1\} \cup \{c_i : i < k_2\}.$$

Suppose $d \in X$. If $f \in (F^B)^n$, define $\mathrm{eval}^D(f, d) = \mathrm{eval}^B(f, d)$ if $d \in B$ and $\mathrm{eval}^D(f, d) = d$ otherwise. Likewise, if $f \in (F^C)^n$ and $d \in C$, put $\mathrm{eval}^D(f, d) = \mathrm{eval}^C(f, c)$ if $c \in C$ and $\mathrm{eval}^C(f, c) = c$ otherwise. If $f \in (F^D)^n \setminus ((F^B)^n \cup (F^C)^n)$, put $\mathrm{eval}^D(f, d) = d$. This defines eval on $(F^D)^n \times X$. More generally, if $f \in (F^D)^n$ and $e \in O^D$, define $\mathrm{eval}^D(f, e) = \mathrm{eval}^D(f, d)$ for the unique $d \in X$ equivalent to e. This is well-defined as B and C agree on A and the D defined in this way is clearly in $\mathbb{K}_n$.

Finally, note that a structure in $\mathbb{K}_n$ generated by k elements is obtained by applying $\leq k^n$ functions of the form $\mathrm{eval}(f, -)$ to $\leq k$ elements in O, so has cardinality $\leq k^{n+1} + k$. This shows $\mathbb{K}_n$ is uniformly locally finite. □

It follows that there is a complete $\aleph_0$-categorical theory T_n^* extending T_n whose models have age $\mathbb{K}_n$ [Hod93, Chapter 7]. By the uniform local finiteness of $\mathbb{K}_n$, T_n^* has quantifier-elimination so T_n^* is the model completion of T_n. Let $\mathbb{M}_n \models T_n^*$ be a monster model.

Definition 3.9.5. Define a ternary relation $\mathop{\smile\!\!\!\!\!\!|}^*$ on small subsets of $\mathbb{M}_n$ by: $a \mathop{\smile\!\!\!\!\!\!|}^*_C b$ if and only if

1. $\mathrm{dcl}(aC)/E \cap \mathrm{dcl}(bC)/E \subseteq \mathrm{dcl}(C)/E$.
2. $\mathrm{dcl}(aC) \cap \mathrm{dcl}(bC) \subseteq \mathrm{dcl}(C)$.

where $X/E = \{[x]_E : x \in X\}$ denotes the collection of E-classes represented by an element of X.

Lemma 3.9.6. *The relation $\mathop{\smile\!\!\!\!\!\!|}^*$ satisfies the independence theorem over structures: if $M \models T_n$ (not necessarily T_n^*), $a \equiv_M a'$, $a \mathop{\smile\!\!\!\!\!\!|}^*_M B$, $a' \mathop{\smile\!\!\!\!\!\!|}^*_M C$ and $B \mathop{\smile\!\!\!\!\!\!|}^*_M C$ then there is a'' with $a'' \equiv_{MB} a$, $a'' \equiv_{MC} a'$, and $a'' \mathop{\smile\!\!\!\!\!\!|}^*_M BC$.*

Proof. We may assume M is a substructure of $\mathbb{M}_n$, $M \subseteq B, C$ and that B and C are definably closed. Write $a = (d_0, \ldots, d_{k-1}, e_0, \ldots, e_{l-1})$ with $d_i \in F$ and $e_j \in O$ and likewise $a' = (d'_0, \ldots, d'_{k-1}, e'_0, \ldots, e'_{l-1})$. Fix an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}_n/M)$ with $\sigma(a) = a'$. Let $U = \{u_f : f \in \mathrm{dcl}(aB) \setminus B\}$ and $V = \{v_f : f \in \mathrm{dcl}(a'C) \setminus C\}$ denote a collection of new formal elements with $u_h = v_{\sigma(h)}$ for all $h \in \langle aM \rangle \setminus B$. Let, then, a_* be defined by

$$a_* = (u_{d_0}, \ldots, u_{d_{k-1}}, u_{e_0}, \ldots, u_{e_{l-1}}) = (v_{d'_0}, \ldots, v_{d'_{k-1}}, v_{e'_0}, \ldots, v_{e'_{l-1}}).$$

We will construct by hand an L-structure D extending $\langle BC \rangle$ with domain $UV\langle BC \rangle$ in which $a^* \equiv_B a$, $a^* \equiv_C a'$ and $a^* \mathop{\smile\!\!\!\!\!\!|}^*_M BC$.

There is a bijection $\iota_0 : \mathrm{dcl}(aB) \to BU$ given by $\iota_0(b) = b$ for all $b \in B$ and $\iota_0(f) = u_f$ for all $f \in \mathrm{dcl}(aB)\backslash B$. Likewise, we have a bijection $\iota_1 : \mathrm{dcl}(a'C) \to CV$ given by $\iota_1(c) = c$ for all

$c \in C$ and $\iota_1(f) = v_f$ for all $f \in \mathrm{dcl}(a'C)\backslash C$. The union of the images of these functions is the domain of the structure D to be constructed and their intersection is $\iota_0(\langle aM\rangle) = \iota_1(\langle a'M\rangle)$. Consider BU and CV as L_n-structures by pushing forward the structure on $\mathrm{dcl}(aB)$ and $\mathrm{dcl}(a'C)$ along ι_0 and ι_1, respectively. Note that $\iota_0|_{\langle aM\rangle} = (\iota_1 \circ \sigma)|_{\langle aM\rangle}$.

We are left to show that we can define an L_n-structure on $UV\langle BC\rangle$ extending that of BU, CV, and $\langle BC\rangle$ in such a way as to obtain a model of T_n^*. To begin, interpret the predicates by $O^D = O^{BU} \cup O^{CV} \cup O^{\langle BC\rangle}$ and $F^D = F^{BU} \cup F^{CV} \cup F^{\langle BC\rangle}$. Let E^D be defined to be the equivalence relation generated by E^{BU}, E^{CV}, and $E^{\langle BC\rangle}$. The interpretation of the predicates is well-defined since if f is an element of $\iota_0(\langle aM\rangle) = \iota_1(\langle a'M\rangle)$ then $\iota_0^{-1}(f)$ is in the predicate O if and only if $\iota_1^{-1}(f)$ is as well, and, moreover, it is easy to check that our assumptions on a, a', B, C entail that no pair of inequivalent elements in BU, CV, or $\langle BC\rangle$ become equivalent in D.

All that is left is to define the function eval^D extending $\mathrm{eval}^{BU} \cup \mathrm{eval}^{CV} \cup \mathrm{eval}^{\langle BC\rangle}$. We first claim that $\mathrm{eval}^{BU} \cup \mathrm{eval}^{CV} \cup \mathrm{eval}^{\langle BC\rangle}$ is a function. The intersection of the domains of the first two functions is (in a Cartesian power of) $\iota_0(\langle aM\rangle) = \iota_1(\langle aM\rangle)$. If b, b' are in this intersection, we must show

$$\mathrm{eval}^{BU}(b, b') = c \iff \mathrm{eval}^{CV}(b, b') = c.$$

Choose $b_0, b_0', c_0 \in \langle aM\rangle$ and $b_1, b_1', c_1 \in \langle a'M\rangle$ with $\iota_i(b_i, b_i', c_i) = (b, b', c)$ for $i = 0, 1$. Then since $\iota_0 = \iota_1 \circ \sigma$ on $\langle aM\rangle$, we have

$$\begin{aligned} \mathbb{M}_n \models \mathrm{eval}(b_0, b_0') = c_0 \quad &\iff \quad \mathbb{M}_n \models \mathrm{eval}(\sigma(b_0), \sigma(b_0')) = \sigma(c_0) \\ &\iff \quad \mathbb{M}_n \models \mathrm{eval}(b_1, b_1') = c_1. \end{aligned}$$

Since eval^{BU} and eval^{CV} are defined by pushing forward the structure on $\langle aB\rangle$ and $\langle a'C\rangle$ along ι_0 and ι_1, respectively, this shows that $\mathrm{eval}^{BU} \cup \mathrm{eval}^{CV}$ defines a function. Now the intersection of $\langle BC\rangle$ with $BU \cup CV$ is BC and, by construction, all 3 functions agree on this set. So the union defines a function.

Choose a complete set of E^D-class representatives $\{d_i : i < \alpha\}$ so that if d_i represents an E^D-class that meets M then $d_i \in M$. If $e \in O^D$ is E^D-equivalent to some e' and (f, e') is in the domain of $\mathrm{eval}^{BU} \cup \mathrm{eval}^{CV} \cup \mathrm{eval}^{\langle BC\rangle}$, define $\mathrm{eval}^D(f, e)$ to be the value that this function takes on (f, e'). On the other hand, if $f \in (F^D)^n \setminus ((F^{BU})^n \cup (F^{CV})^n \cup (F^{\langle BC\rangle})^n)$ or e is not E^D-equivalent to any element on which $\mathrm{eval}^D(f, -)$ has already been defined, put $\mathrm{eval}^D(f, e) = d_i$ for the unique d_i which is E^D-equivalent to e. This now defines eval^D on all of $(F^D)^n \times O^D$ and, by construction, $\mathrm{eval}^D(f, -)$ is a selector function for E^D for all $f \in (F^D)^n$. This completes the construction of D and we have shown D is a model of T_n. By model-completeness and saturation, D embeds into $\mathbb{M}_n$ over BC. If we can show $a_* \mathop{\smile\!\!\!\!\!\!\!\;|}\nolimits^*_M BC$ in D, then this will be true for the image of a_* in D.

We have already argued that BU and CV are substructures of D—it follows that every E^D-class represented by an element of a_* can only be equivalent to an element of B or C if it is equivalent to an element of M. Moreover, our construction has guaranteed that $\langle a_*M\rangle^D \cap \langle BC\rangle \subseteq BU \cap \langle BC\rangle^D \subseteq B$ and, by similar reasoning, $\langle a_*M\rangle \subseteq C$. This implies $\langle a_*M\rangle^D \cap \langle BC\rangle B \cap C \subseteq M$, so $a_* \mathop{\smile\!\!\!\!\!\!\!\;|}\nolimits^*_M BC$. $\square$

Proposition 3.9.7. *The theory T_n^* is NSOP$_1$ and, moreover, if $M \models T_n^*$, then $a \mathop{\smile\!\!\!\!\!\mid}^*_M b$ if and only if $a \mathop{\smile\!\!\!\!\!\mid}^K_M b$.*

Proof. In Lemma 3.9.6, we showed $\mathop{\smile\!\!\!\!\!\mid}^*$ satisfies the independence theorem over a model, and the other conditions (1)-(4) in Theorem 6.2.4 are clear for $\mathop{\smile\!\!\!\!\!\mid}^*$. To show (6), notice that if $A \not\mathop{\smile\!\!\!\!\!\mid}^*_M B$ with A, B definably closed and containing M, then either there is some $a \in A$ and $b \in B$ so that $\models E(a,b)$ and the E-class of b does not meet M or $a = b$ for some $b \notin M$. Suppose $(b_i)_{i<\omega}$ is a Morley sequence in some global M-invariant $q \supseteq \text{tp}(b/M)$. If the class of b does not meet M, then $\neg E(b_i, b_j)$ for $i \neq j$ by M-invariance so $\{E(x;b_i) : i < \omega\}$ is 2-inconsistent. Likewise, if b is not in M, then $b_i \neq b_j$ for $i \neq j$ so $\{x = b_i : i < \omega\}$ are 2-inconsistent. It follows that $\mathop{\smile\!\!\!\!\!\mid}^* = \mathop{\smile\!\!\!\!\!\mid}^K$ over models. □

We note that $\mathop{\smile\!\!\!\!\!\mid}^K$ satisfies (a form of) local character in T_1^*:

Proposition 3.9.8. *For any model $N \models T_1^*$ and $p \in S(N)$, there is a countable $M \prec N$ so that $p(x)$ does not Kim-fork over M.*

Proof. We use the characterization of $\mathop{\smile\!\!\!\!\!\mid}^K$ from Proposition 3.9.7. Let $a \models p$ and choose $M_0 \prec N$ to be an arbitrary elementary submodel. By induction, construct an elementary chain $(M_i)_{i<\omega}$ of countable elementary submodels of N so that $\text{dcl}(aM_i) \cap N \subseteq M_{i+1}$ and every equivalence class in $\text{dcl}(aM_i)/E \cap N/E$ is represented by an element of M_{i+1}. Since M_i is countable, $\text{dcl}(aM_i)$ is countable, there is no problem in choosing M_{i+1}, by downward Löwenheim-Skolem. Let $M = \bigcup_{i<\omega} M_i$. We claim $a \mathop{\smile\!\!\!\!\!\mid}^K_M N$. Given $c \in \text{dcl}(aM) \cap N$, there is n so that $c \in \text{dcl}(aM_n) \cap N$ hence $c \in M_n \subseteq M$. This shows $\text{dcl}(aM) \cap N \subseteq M$. Arguing similiarly, we have $\text{dcl}(aM)/E \cap N/E \subseteq M/E$. This shows $a \mathop{\smile\!\!\!\!\!\mid}^K_M N$. □

Lemma 3.9.9. *Modulo T_1^*, the formula $O(x)$ axiomatizes a complete type over $\emptyset$ which is not co-simple.*

Proof. That $O(x)$ implies a complete type is clear from quantifier-elimination. In $O(\mathbb{M}_1)$, choose an array $(a_{\alpha,\beta})_{\alpha,\beta<\omega}$ of distinct elements so that, for all $\alpha < \alpha' < \omega$, given β, β', $\mathbb{M}_1 \models E(a_{\alpha,\beta}, a_{\alpha,\beta'})$ and $\mathbb{M}_1 \models \neg E(a_{\alpha,\beta}, a_{\alpha',\beta'})$. Let $\varphi(x;y)$ be the formula $\text{eval}(x,y) = y$. It is now easy to check

- For all functions $f : \omega \to \omega$, $\{\varphi(x; a_{\alpha,f(\alpha)}) : \alpha < \omega\}$ is consistent
- For all $\alpha < \omega$, $\{\varphi(x; a_{\alpha,\beta}) : \beta < \omega\}$ is 2-inconsistent,

so $\varphi(x;y)$ witnesses TP$_2$ with respect to parameters realizing $O(x)$. This shows $O(x)$ is not co-simple. □

Lemma 3.9.10. *Suppose $A \subseteq \mathbb{M}_1$. Then $\text{acl}(A) = \langle A \rangle = A \cup \text{eval}(F(A) \times O(A))$.*

Proof. The equality of $\text{acl}(A)$ and $\langle A \rangle$ follows from SAP for $\mathbb{K}_1$ [Hod93, Theorem 7.1.8]. The axioms of T_1^* imply that every term of L_1 is equivalent to one of the form x or $\text{eval}(x,y)$, so $\langle A \rangle = A \cup \text{eval}(F(A) \times O(A))$. □

We will see that $\mathop{\smile\!\!\!\!\!\!|}^*$ characterizes dividing when elements on the left-hand side come from O. The following lemma is the key ingredient in proving this:

Lemma 3.9.11. *Suppose $A = dcl(A) \subseteq \mathbb{M}_1$ and $A = \langle a, B\rangle$ for some $a \in O(A)$ and $B = dcl(B) \subseteq \mathbb{M}_1$, where $l(a) = 1$. Given a sequence $(B_i)_{i<N}$ of substructures of $\mathbb{M}_1$ isomorphic to B over $C = dcl(C)$ where for $i \neq j$, $B_i \cap B_j = C$. Then if $a \mathop{\smile\!\!\!\!\!\!|}^*_C B$, then there is a structure $D \models T_1$ and some $a' \in D$ so that*

1. $\langle (B_i)_{i<N}\rangle \subseteq D$.
2. $\langle a', B_i\rangle^D \cong_C \langle a, B\rangle$ *for all* $i < N$.

Proof. Suppose $A = \langle a, B\rangle$, $(B_i)_{i<N}$ and C are given as in statement, satisfying (1). If $a \in C$, the lemma is clear so assume it is not, and therefore $a \notin B$ by our assumption that $A \models a \neq b$ for all $b \in B \setminus C$. Moreover, we may assume $B_0 = B$. Note that the underlying set of A is $B \cup \{a\} \cup \mathrm{eval}(F(B), a)$. Let $X = \langle (B_i)_{i<N}\rangle$.

Case 1: $A \models E(a, c)$ for some $c \in C$. In this case, the underlying set of A is $B \cup \{a\} \cup \mathrm{eval}(F(B), c) = B \cup \{a\}$. Let D be the extension of X with underlying set $X \cup \{a\}$ with relations interpreted so that $D \models a \in O \wedge E(a, c)$ and the function eval defined to extend eval^X and so that $\mathrm{eval}^D(d, a) = \mathrm{eval}^X(d, c)$ for all $d \in O^D$. It is easy to check that this satisfies (2).

Case 2: $A \models \neg E(a, c)$ for all $c \in C$. By our assumption that A satisfies (1), it follows that $A \models \neg E(a, b)$ for all $b \in B$ and hence the underlying set of A is the disjoint union of B and $\{a\} \cup \mathrm{eval}^A(F(B), a)$. Let $Y = \{a\} \cup \mathrm{eval}^A(F(B), a)$. We will define an L_1-structure extending X with underlying set $X \cup Y$. Interpret the sorts $F^D = F^X$ and $O^D = O^X \cup Y$. Define the equivalence relation so that $E^X \subset E^D$ and Y forms one E^D-class.

Fix for all $i < N$ a C-isomorphism $\sigma_i : B_i \to B_0$ (assume $\sigma_0 = \mathrm{id}_{B_0}$). Note that $F^X = \bigcup_{i<N} F^{B_i}$. Interpret eval^D to extend eval^X and so that, if $b \in F^{B_i}$ and $e \in Y$,

$$\mathrm{eval}^D(b, e) = \mathrm{eval}^A(\sigma_i(b), a).$$

This defines $D \models T_1$ and, by construction, the map extending σ_i and sending $a \mapsto a$ induces an isomorphism $\langle a, B_i\rangle^D \to \langle a, B_0\rangle^D = A$ for all $i < N$. This completes the proof. □

Corollary 3.9.12. *Suppose F is a substructure of $\mathbb{M}_1$. If $a \in O(\mathbb{M}_1)$ and $l(a) = 1$, then $a \mathop{\smile\!\!\!\!\!\!|}^d_F B$ if and only if $a \mathop{\smile\!\!\!\!\!\!|}^*_F B$.*

Proof. If $a \not\mathop{\smile\!\!\!\!\!\!|}^*_F B$ then clearly $a \not\mathop{\smile\!\!\!\!\!\!|}^d_F B$, so we prove the other direction. Suppose $a \mathop{\smile\!\!\!\!\!\!|}^*_F B$ and $a \not\mathop{\smile\!\!\!\!\!\!|}^d_F B$ and we will get a contradiction. Suppose $\varphi(x; c, b)$ witnesses dividing, so $\varphi(x; c, b) \in \mathrm{tp}(a/FB)$ with $c \in F$ and $b \in B$, and there is an F-indiscernible sequence $\langle b_i : i < \omega\rangle$ with $b_0 = b$ so that $\{\varphi(x; c, b_i) : i < \omega\}$ is k-inconsistent for some k. As $a \mathop{\smile\!\!\!\!\!\!|}^*_F B$ we may, by growing the tuple c, assume that every equivalence class represented both by b and a is represented by an element of c. Let $B_i = \langle c, b_i\rangle$, $C = B_0 \cap B_i$ for some/all $i \neq 0$ (by F-indiscernibility, this is well-defined and contains c) and $A = \langle a, b, c\rangle$. As $a \mathop{\smile\!\!\!\!\!\!|}^*_C B$, the structures A, C, and $(B_i)_{i<k+1}$ satisfy (1) of Lemma 6.3.17, and therefore there is $D \in \mathbb{K}_1$ and some $a' \in D$ so that $\langle (B_i)_{i<k+1} \rangle \subseteq D$ and $\langle a', B_i \rangle^D \cong \langle a, B \rangle^A$ for all $i < k+1$. By embedding D into $\mathbb{M}_1$ over $\langle (B_i)_{i<k+1} \rangle^D$ we see that, in $\mathbb{M}_1$, $\{\varphi(x; c, b_i) : i < k+1\}$ is consistent by quantifier-elimination. This is a contradiction. □

Corollary 3.9.13. *The theory T_1^* is NSOP$_1$ and the formula $O(x)$ axiomatizes a complete type which is simple and not cosimple.*

Proof. Lemma 3.9.9 shows that $O(x)$ axiomatizes a complete type which is not cosimple. To show $O(x)$ is simple, we have to show that $\mathop{\smile\!\!\!\!\!\!\mid}^d$ satisfies local character on $O(x)$. So fix any $a \in \mathbb{M}_1$ with $\mathbb{M}_1 \models O(a)$ and any small set $B \subseteq \mathbb{M}_1$. We may suppose $B = \text{dcl}(B)$. Notice that $\text{dcl}(a) = a$. If $a \in B$ then $a \mathop{\smile\!\!\!\!\!\!\mid}^*_a B$. If $a \notin B$ but $\mathbb{M} \models E(a, b)$ for some $b \in B$ then $a \mathop{\smile\!\!\!\!\!\!\mid}^*_b B$. Finally, if a not E-equivalent to any element of B then $a \mathop{\smile\!\!\!\!\!\!\mid}^*_\emptyset B$. Corollary 3.9.12 showed $a \mathop{\smile\!\!\!\!\!\!\mid}^*_C B$ if and only if $a \mathop{\smile\!\!\!\!\!\!\mid}^d_C B$ for any a with $\mathbb{M} \models O(a)$, so $\mathop{\smile\!\!\!\!\!\!\mid}^d$ satisfies local character on O. Therefore O is simple. □

Remark 3.9.14. This answers Problem 6.10 of [Che14].

Remark 3.9.15. Given a model $M \models T_1^*$, one can consider the complete type $p(x)$ over M axiomatized by saying

- $O(x)$
- $\neg E(x, m)$ for all $m \in O(M)$
- $\text{eval}(m, x) \neq x$ for all $m \in O(M)$

In a similar fashion, one can check that this is simple, non-co-simple so, in particular, nothing is gained by working over a model. In fact, in this situation, we get another proof of the corollary, using Proposition 3.8.3, as we have shown that if $a \models p$, then $a \mathop{\smile\!\!\!\!\!\!\mid}^d_M b$ if and only if $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ so p is simple.

Proposition 3.8.7 above shows that in any non-simple NSOP$_1$ theory, there are types over models with no universal Morley sequences in them. The following gives an explicit example:

Proposition 3.9.16. *Given $M \models T_1^*$, there is a type $p \in S(M)$ with no universal Morley sequence.*

Proof. Pick $b \in O(\mathbb{M})$ not in M and let $p(x) = \text{tp}(b/M)$. Towards contradiction, suppose $(b_i)_{i<\omega}$ is a universal Morley sequence in p.

Case 1: $\mathbb{M} \models E(b_i, b_j)$ for all $i, j < \omega$.

The formula $E(x; b)$ divides over M: choose any M-indiscernible sequence $\langle c_i : i < \omega \rangle$ with $c_0 = b$ and $\neg E(c_i, c_{i+1})$ – then $\{E(x; c_i) : i < \omega\}$ is inconsistent. But $\{E(x; b_i) : i < \omega\}$ is consistent, a contradiction.

Case 2: $\mathbb{M} \models \neg E(b_i, b_j)$ for $i \neq j$. The formula $\text{eval}(x, b) = b$ divides over M – choose any M-indiscernible sequence $\langle c_i : i < \omega \rangle$ with $E(c_i, c_j)$ for all i, j and $c_0 = b$. Then

$\{\mathrm{eval}(x, c_i) = c_i : i < \omega\}$ is inconsistent (as for any a, the function $\mathrm{eval}(a, -)$ takes on only one value on elements of any equivalence class). But $\{\mathrm{eval}(x, b_i) = b_i : i < \omega\}$ is consistent, a contradiction. □

Proposition 3.9.17. *In T_1^*, forking does not equal dividing, even over models.*

Proof. Fix $M \models T_1^*$. Let $\varphi(x, y; z)$ be the formula $\mathrm{eval}(x, z) = z \vee E(y, z)$. Given any $b \in O(\mathbb{M}_1)$ not in M, we claim the formula $\varphi(x, y; b)$ forks but does not divide over M. The proof of Proposition 3.9.16 shows that both $E(x, b)$ and $\mathrm{eval}(x, b) = b$ divide over M so $\varphi(x, y; b)$ forks over M. Given any M-indiscernible sequence $\langle b_i : i < \omega\rangle$ starting with b, either all b_i's lie in a single equivalence class, in which case $\{E(y, b_i) : i < \omega\}$ is consistent, or they all lie in different classes, in which case $\{\mathrm{eval}(x, b_i) = b_i : i < \omega\}$ is consistent. Either way, $\{\varphi(x, y; b_i) : i < \omega\}$ is consistent, so $\varphi(x, y; b)$ does not divide over M. □

Lemma 3.9.18. *Any a be a tuple in F, b a tuple in O, and $C = dcl(C) \subseteq \mathbb{M}$. Then $tp(a, b/C)$ extends to a global C-invariant type.*

Proof. Write $a = (a_0, \dots, a_{n-1})$, $b = (b_0, \dots, b_{k-1})$. We may assume that no equalities occur between the elements of a and of b, or between a, b and C. We define a C-invariant global type $p(x, y) \in S(\mathbb{M})$ as follows. The type $p(x, y)$ contains all formulas of $\mathrm{tp}(a, b/C)$ together with the following axiom scheme:

$$\begin{aligned}
\mathrm{eval}(x_i, m) \neq m &\in p(x, y) \text{ for all } i < n, m \in \mathbb{M} \setminus C.\\
\mathrm{eval}(x_i, m) \neq \mathrm{eval}(x_j, m) &\in p(x, y) \text{ for all } i < j < n, m \in \mathbb{M} \text{ with } m/E \notin C/E.\\
\mathrm{eval}(x_i, y_j) \neq m &\in p(x, y) \text{ for all } i < n, j < k, m \in \mathbb{M} \setminus C.\\
\mathrm{eval}(m, y_j) \neq y_j &\in p(x, y) \text{ for all } j < k, m \in \mathbb{M} \setminus C.\\
\neg E(y_j, m) &\in p(x, y) \text{ for all } j < k, m \in \mathbb{M} \text{ with } m/E \notin C/E.
\end{aligned}$$

It is clear that this type is consistent and C-invariant. We claim it implies a complete type over $\mathbb{M}$: note that because $\mathrm{eval}(x, \mathrm{eval}(y, z)) = \mathrm{eval}(x, z)$, every term is equivalent to x or $\mathrm{eval}(x, y)$. Because $E(x, \mathrm{eval}(y, z))$ is equivalent to $E(x, z)$, every atomic formula is equivalent to an equality of terms or of the form $E(x, y)$. Equalities of the form $\mathrm{eval}(x_i, y_j) = \mathrm{eval}(x_{i'}, y_{j'})$ are implied or negated by $\mathrm{tp}(a, b/C)$, so the truth value of every atomic formula in the variables x, y with parameters in $\mathbb{M}$ is determined by the above. □

Corollary 3.9.19. *The theory T_1^* is an $NSOP_1$ theory for which forking does not equal dividing, yet every type has a global non-forking extension.*

Remark 3.9.20. This answers Question 7.1(1) of [Con14], which asked if forking = dividing in every NSOP_3 theory in which every type has a global non-forking extension, as every NSOP_1 theory is NSOP_3 [DS04, Claim 2.3].

Finally, the following proposition gives a counter-example to the form of transitivity mentioned at the beginning of the subsection.

Proposition 3.9.21. *For any model $M \models T_2^*$, there are f, g, and c so that $f \mathop{\smile\!\!\!\!\!\!\;|}^K_M gc$, $g \mathop{\smile\!\!\!\!\!\!\;|}^K_M c$, and $fg \not\mathop{\smile\!\!\!\!\!\!\;|}^K_M c$.*

Proof. Given $M \models T_2^*$, choose any $c \in \mathbb{M}_2 \setminus M$ in an E-class represented by an element m of M—let $\{m_i : i < \alpha\}$ enumerate a set of representatives for the remaining E-classes of M. Then choose distinct elements $f, g \in F$ so that

1. $\text{eval}(f, g, m) = \text{eval}(g, f, m) = c$.
2. $\text{eval}(f, h, m) = \text{eval}(h, f, m) = m$ and
$$\text{eval}(f, h, m_i) = \text{eval}(h, f, m_i) = m_i$$
for all $h \in F^M \cup \{f\}$.
3. $\text{eval}(g, h, m) = \text{eval}(h, g, m) = m$ and
$$\text{eval}(g, h, m_i) = \text{eval}(h, g, m_i) = m_i$$
for all $h \in F^M \cup \{g\}$.

Then we have

$$\begin{array}{rcl} \text{dcl}(fM) & = & M \cup \{f\} \\ \text{dcl}(gM) & = & M \cup \{g\} \\ \text{dcl}(cM) & = & M \cup \{c\} \\ \text{dcl}(fgM) & = & M \cup \{f, g, c\} \\ \text{dcl}(gcM) & = & M \cup \{g, c\}. \end{array}$$

It follows that $\text{dcl}(fM) \cap \text{dcl}(gcM)$ and $\text{dcl}(gM) \cap \text{dcl}(cM)$ are contained in M so $f \mathop{\smile\!\!\!\!\!\!\;|}^*_M gc$ and $g \mathop{\smile\!\!\!\!\!\!\;|}^*_M c$. However, $c \in (\text{dcl}(fgM) \cap \text{dcl}(cM)) \setminus M$, showing $fg \not\mathop{\smile\!\!\!\!\!\!\;|}^*_M c$. As Proposition 3.9.7 showed $\mathop{\smile\!\!\!\!\!\!\;|}^K = \mathop{\smile\!\!\!\!\!\!\;|}^*$, we are done. $\square$

Frobenius Fields

In this section, we study a class of NSOP$_1$ fields. If F is a field, we write F^{alg} and F^s for the algebraic and separable closures of F, respectively.

Definition 3.9.22. Suppose F is a field.

1. We say F is *pseudo-algebraically closed* (PAC) if every absolutely irreducible variety over F has an F-rational point.

2. We say F is a *Frobenius field* if F is PAC and its absolute Galois group $\mathcal{G}(F)$ has the *embedding property* (also known as the *Iwasawa property*), that is, if $\alpha : \mathcal{G}(F) \to A$ and $\beta : B \to A$ are continuous epimorphisms and B is a finite quotient of $\mathcal{G}(F)$, then there is a continuous epimorphism $\gamma : \mathcal{G}(F) \to B$ so that $\beta \circ \gamma = \alpha$ as in the following diagram:

$$\begin{tikzcd} & \mathcal{G}(F) \arrow[dl, dashed, two heads] \arrow[d, two heads] \\ B \arrow[r, two heads] & A \end{tikzcd}$$

The free profinite group on countably many generators $\hat{F}_\omega$ has the embedding property so the ω-free PAC fields are Frobenius fields. However, there are many others—see, e.g., [FJ08, 24.6].

Definition 3.9.23. Suppose G is a profinite group. Let $\mathcal{N}(G)$ be the collection of open normal subgroups of G. We define

$$\mathcal{S}(G) = \coprod_{N \in \mathcal{N}(G)} G/N.$$

Let L_G the language with a sort X_n for each $n \in \mathbb{Z}^+$, two binary relation symbols $\leq$, C, and a ternary relation P. We regard $\mathcal{S}(G)$ as an L_G-structure in the following way:

- The coset gN is in sort X_n if and only if $[G : N] \leq n$.
- $gN \leq hM$ if and only if $N \subseteq M$
- $C(gN, hM) \iff N \subseteq M$ and $gM = hM$.
- $P(g_1N_1, g_2N_2, g_3N_3) \iff N_1 = N_2 = N_3$ and $g_1g_2N_1 = g_3N_1$.

Note that we do not require that the sorts be disjoint (see [Cha98, Section 1] for a discussion on the syntax of this structure).

Interpretability of $\mathcal{S}(\mathcal{G}(F))$ in (F^{alg}, F) is proved in [Cha02, Proposition 5.5]. The "moreover" clause is clear from the proof.

Fact 3.9.24. Both F and $\mathcal{S}(\mathcal{G}(F))$ are interpretable in (K, F) where K is any algebraically closed field containing F. Call the interpretation π. Moreover, if $L \subseteq F$ is a subfield so that F is a regular extension of L, then the restriction of π to (K, L) produces an interpretation of $\mathcal{S}(\mathcal{G}(L))$, contained in $\mathcal{S}(\mathcal{G}(F))$ in a natural way.

Lemma 3.9.25. *Let F be a large sufficiently saturated and homogeneous field (i.e. a monster model of its theory) and $M \prec F$ a small elementary substructure. Suppose $A = acl(A)$, $B = acl(B)$ are subsets of F with $M \subseteq A \cap B$.*

1. *If $A \equiv_M B$ in F, then $\mathcal{S}(\mathcal{G}(A)) \equiv_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(B))$.*

2. *If $(A_i)_{i<\omega}$ is an M-indiscernible sequence with $A_0 = A$, then $(\mathcal{S}(\mathcal{G}(A_i)))_{i<\omega}$ is $\mathcal{S}(\mathcal{G}(M))$-indiscernible.*

3. *If $A \mathop{\smile\!\!\!\!\!\!\mid}^u_M B$ in F, then $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!\mid}^u_{\mathcal{S}(\mathcal{G}(B))} \mathcal{S}(\mathcal{G}(B))$ in $\mathcal{S}(\mathcal{G}(M))$.*

Proof. (1) If $A \equiv_M B$ in F, then there is an automorphism $\sigma \in \mathrm{Aut}(F/M)$ with $\sigma(A) = B$. The map σ has an extension $\tilde{\sigma}$ to F^{alg} which is, then, an automorphism of the pair (F^{alg}, F) taking A to B and fixing M pointwise. It follows $A \equiv_M B$ in the pair (F^{alg}, F). Since $A = \mathrm{acl}(A)$ and $B = \mathrm{acl}(B)$, we know F is a regular extension of A and of B (see, e.g., [Cha99, Section 1.17]). By Fact 7.1.7, we have $\mathcal{S}(\mathcal{G}(A)) \equiv_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(B))$.

(2) If $(A_i)_{i<\omega}$ is an M-indiscernible sequence with $A_0 = A$, given $i_0 < \ldots < i_{k-1}$ and $j_0 < \ldots < j_{k-1}$, we know $A_{i_0} \ldots A_{i_{k-1}} \equiv_M A_{j_0} \ldots A_{j_{k-1}}$ so $\mathrm{acl}(A_{i_0} \ldots A_{i_{k-1}}) \equiv_M \mathrm{acl}(A_{j_0} \ldots A_{j_{k-1}})$. Then by (1) $\mathcal{S}(\mathcal{G}(\mathrm{acl}(A_{i_0} \ldots A_{i_{k-1}}))) \equiv_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(\mathrm{acl}(A_{j_0} \ldots A_{j_{k-1}})))$, which implies $(\mathcal{S}(\mathcal{G}(A_i)))_{i<\omega}$ is $\mathcal{S}(\mathcal{G}(M))$-indiscernible.

(3) In *any* theory, if π is an interpretation of the structure X in the structure Y, and $A \mathop{\smile\!\!\!\!\!\!\mid}^u_C B$ in Y, then $\pi(A) \mathop{\smile\!\!\!\!\!\!\mid}^u_{\pi(C)} \pi(B)$. It follows that if $A \mathop{\smile\!\!\!\!\!\!\mid}^u_M B$ in F, then $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!\mid}^u_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(\mathcal{B}))$ by Fact 7.1.7. □

Proposition 3.9.26. *Suppose F is an arbitrary field and, in an elementary extension F^* of F, $a \mathop{\smile\!\!\!\!\!\!\mid}^K_F b$. Then the fields $A = acl(Fa)$ and $B = acl(Fb)$ satisfy the following conditions:*

1. *A and B are linearly disjoint over F*

2. *F^* is a separable extension of AB*

3. *$acl(AB) \cap A^s B^s = AB$.*

Proof. In [Cha99, Theorem 3.5], Chatizdakis proves (1)-(3) for an *arbitrary* theory of fields under the assumption that $a \mathop{\smile\!\!\!\!\!\!\mid}^f_F b$. She deduces from $a \mathop{\smile\!\!\!\!\!\!\mid}^f_F b$ that there is an F-indiscernible coheir sequence $(B_i)_{i<\omega}$, i.e. an F-indiscernible sequence with $B_{<i} \mathop{\smile\!\!\!\!\!\!\mid}^u_F B_i$ for all i, so that $AB_i \equiv_F AB$ for all i (rather, she proves this with a *heir* sequence, but the argument is symmetric). She then proves that (1)-(3) follow from the existence of such a sequence. Note, however, that this follows merely from the assumption $a \mathop{\smile\!\!\!\!\!\!\mid}^K_F b$. □

Remark 3.9.27. Note (1) and (2) are equivalent to saying $A \mathop{\smile\!\!\!\!\!\!\mid}^{SCF}_F B$ [Cha99, Remark 3.3], where SCF denotes the complete (stable) theory of which F^s is a model.

Lemma 3.9.28. *Suppose F is a Frobenius field. If $A = acl(A)$, $B = acl(B)$ contain F and $A \mathop{\smile\!\!\!\!\!\!\mid}^K_F B$ then $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!\mid}^f_{\mathcal{S}(\mathcal{G}(F))} \mathcal{S}(\mathcal{G}(B))$ in $Th(\mathcal{S}(\mathcal{G}(F)))$.*

Proof. Chatzidakis [Cha98] shows that the Galois group $\mathcal{S}(\mathcal{G}(F))$ is ω-stable. Let $(B_i)_{i<\omega}$ be a Morley sequence in a global type finitely satisfiable in F extending $\mathrm{tp}(B/F)$. As $A \mathop{\smile\!\!\!\!\!\!\mid}^K_F B$, we may assume $(B_i)_{i<\omega}$ is A-indiscernible. Then $(\mathcal{S}(\mathcal{G}(B_i)))_{i<\omega}$ is a Morley sequence in a global type finitely satisfiable in $\mathcal{S}(\mathcal{G}(F))$ which is moreover $\mathcal{S}(\mathcal{G}(A))$-indiscernible. This implies $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!\mid}^K_{\mathcal{S}(\mathcal{G}(F))} \mathcal{S}(\mathcal{G}(B))$. As $\mathrm{Th}(\mathcal{S}(\mathcal{G}(F))$ is simple, this implies $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!\mid}^f_{\mathcal{S}(\mathcal{G}(F))} \mathcal{S}(\mathcal{G}(B))$ by Kim's lemma [Kim98, Proposition 2.1]. □

Fix a field F and let SCF denote the complete theory of which F^s is a model.

Definition 3.9.29. Suppose $A = \mathrm{acl}(A)$, $B = \mathrm{acl}(B)$, and $C = \mathrm{acl}(C)$ in the field F. We say A is *weakly independent* from B over C if

1. $A \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{SCF}}_C B$
2. $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!\mid}^f_{\mathcal{S}(\mathcal{G}(F))} \mathcal{S}(\mathcal{G}(B))$, where $\mathop{\smile\!\!\!\!\!\!\mid}^f$ denotes non-forking independence in $\mathrm{Th}(\mathcal{S}(\mathcal{G}(\mathcal{F})))$

Extend this to arbitrary tuples by stipulating a is *weakly independent* from b over c if and only if $\mathrm{acl}(a, c)$ is weakly independent from $\mathrm{acl}(b, c)$ over $\mathrm{acl}(c)$.

Theorem 3.9.30. *[Cha02, Theorem 6.1] Let F be a Frobenius field, sufficiently saturated, and $E = acl(E)$ a subfield of F. Assume, moreover, that $acl(\mathcal{S}(\mathcal{G}(E))) = \mathcal{S}(\mathcal{G}(E))$ and if the degree of imperfection of F is finite, that E contains a p-basis of F. Assume that the tuples a, b, c_1, c_2 of F satisfy:*

1. *a and c_1 are weakly independent over E, b and c_2 are weakly independent over E, $c_1 \equiv_E c_2$*
2. *$acl(Ea)$ and $acl(Eb)$ are SCF-independent over E.*

Then there is c realizing $tp(acl(Ea)) \cup tp(c_2/acl(Eb))$ such that c and $acl(Eab)$ are weakly independent over E.

Theorem 3.9.31. *Suppose F is a Frobenius field and a, b are tuples from an elementary extension of F. Then $a \mathop{\smile\!\!\!\!\!\!\mid}^K_F b$ if and only if a and b are weakly independent over F.*

Proof. Given a, b, and F, set $A = \mathrm{acl}(aF)$ and $B = \mathrm{acl}(bF)$. It suffices to show $A \mathop{\smile\!\!\!\!\!\!\mid}^K_F B$ if and only if A is weakly independent from B over F. If $A \mathop{\smile\!\!\!\!\!\!\mid}^K_F B$, then $A \mathop{\smile\!\!\!\!\!\!\mid}^{SCF}_F B$ by Proposition 3.9.26 and $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!\mid}^f_{\mathcal{S}(\mathcal{G}(F))} \mathcal{S}(\mathcal{G}(B))$ by Proposition 3.9.26. Hence A and B are weakly independent over F. For the other direction, suppose A and B are weakly independent over F. Let $(B_i)_{i<\omega}$ be a Morley sequence in a global F-invariant type with $B_0 = B$ and set $p(X; B) = \mathrm{tp}(A/B)$. We will show by induction that $\bigcup_{i \leq n} p(X; B_i)$ has a realization weakly independent from $(B_i)_{i \leq n}$ over F. For $n = 0$, this is by the assumption that A and B are weakly independent over F. If it has been shown for n, then note that, because, $B_{n+1} \mathop{\smile\!\!\!\!\!\!\mid}^i_F B_0 \dots B_n$, we have, in particular, B_{n+1} and $(B_i)_{i \leq n}$ are weakly independent over F. By Theorem 3.9.30, $p(X; B_{n+1}) \cup \bigcup_{i \leq n} p(X; B_i)$ has a realization weakly independent from $(B_i)_{i \leq n+1}$. By compactness, we conclude $\bigcup_{i<\omega} p(X; B_i)$ is consistent. As $(B_i)_{i<\omega}$ was arbitrary, this shows $A \mathop{\smile\!\!\!\!\!\!\mid}^K_F B$. □

Vector spaces

The theories of a vector space over a field equipped with a symmetric or alternating bilinear form have model companions—they are the theories of an infinite dimensional vector space over an algebraically closed field equipped with a generic nondegenerate alternating or symmetric bilinear form. We use T_∞ to refer to both the model companion where the form is symmetric and where it is alternating, as this choice makes no difference for our analysis below. The language is two-sorted: there is a sort V for the vector space, with the language of abelian groups on it, a sort K for the field, equipped with the ring language, a function $K \times V \to V$ for the action of scalar multiplication, and a function $[,] : V \times V \to K$ for the bilinear form. In this subsection, we write $\mathbb{M} \models T_\infty$ for a fixed monster model of T_∞.

Fact 3.9.32. Given a set $X \subseteq \mathbb{M}$, write X_K for the field points of X and X_V for the vector space points of X. For Y a set of vectors, write $\langle Y \rangle$ for the $\mathbb{M}_K$-span of V.

1. T_∞ eliminates quantifiers after expanding the vector space sort with an n-ary predicate θ_n interpreted so that $\models \theta_n(v_0, \ldots, v_{n-1})$ if and only if $v_0, \ldots, v_{n-1}$ are linearly independent for all $n \geq 2$ [Gra99, Theorem 9.2.3].

2. For any set $A \subseteq \mathbb{M}$, the field points of $\mathrm{dcl}(A)$ contain the field generated by A_K, $\{[a, b] : a, b \in A_V\}$, and for each n, and every set $\{\alpha_0, \ldots, \alpha_{n-1}\}$ such that there are $v_0, \ldots, v_n \in A_V$ with $\mathbb{M} \models \theta_n(v_0, \ldots, v_{n-1})$ and $v_n = \alpha_0 v_0 + \ldots + \alpha_{n-1} v_{n-1}$. The vector space points of $\mathrm{dcl}(A)$ are the $(\mathrm{dcl}(A))_K$-span of A_V. The field points of $\mathrm{acl}(A)$ are the algebraic closure of $(\mathrm{dcl}(A))_K$ and the vector space points of $\mathrm{acl}(A)$ are the $(\mathrm{acl}(A))_K$-span of A_V [Gra99, Proposition 9.5.1].

Definition 3.9.33. Write $\mathop{\smile\!\!\!\!\!\!|}^{ACF}$ to denote algebraic independence, which coincides with non-forking independence in the theory ACF. Suppose $A \subseteq B$ and c is a singleton. Let $c \mathop{\smile\!\!\!\!\!\!|}^{\Gamma}_A B$ be the assertion that $(\mathrm{dcl}(cA))_K \mathop{\smile\!\!\!\!\!\!|}^{ACF}_{(\mathrm{dcl}(A))_K} (\mathrm{dcl}(B))_K$ and one of the following holds:

1. $c \in \mathbb{M}_K$

2. $c \in \langle A_V \rangle$

3. $c \notin \langle B_V \rangle$ and $[c, B]$ is Φ-independent over A,

where '$[c, B]$ is Φ-independent over A' means that whenever $\{b_0, \ldots, b_{n-1}\}$ is a linearly independent set in $B_V \cap (\mathbb{M}_V \setminus \langle A \rangle)$ then the set $\{[c, b_0], \ldots, [c, b_{n-1}]\}$ is algebraically independent over the compositum of $(\mathrm{dcl}(B))_K$ and $(\mathrm{dcl}(Ac))_K$.

By induction, for $c = (c_0, \ldots, c_m)$ define $c \mathop{\smile\!\!\!\!\!\!|}^{\Gamma}_A B$ by

$$c \mathop{\smile\!\!\!\!\!\!|}^{\Gamma}_A B \iff (c_0, \ldots, c_{m-1}) \mathop{\smile\!\!\!\!\!\!|}^{\Gamma}_A B \text{ and } c_m \mathop{\smile\!\!\!\!\!\!|}^{\Gamma}_{Ac_0 \ldots c_{m-1}} Bc_0 \ldots c_{m-1}.$$

Fact 3.9.34. [Gra99, Theorem 12.2.2] The relation $\mathop{\smile\!\!\!\!\!\!|}^{\Gamma}$ is automorphism invariant and symmetric. Moreover, it satisfies extension, strong finite character, and the independence theorem over a model. Consequently, T_∞ is NSOP$_1$.

Proposition 3.9.35. *Suppose $M \models T_\infty$. Then if $A = acl(A)$, $B = acl(B)$ and $A \cap B \supseteq M$, then $A \mathop{\smile\!\!\!\!\!\!|}^{K}_{M} B$ if and only if $A \cap B = M$.*

Proof. The right to left direction is trivial and holds in any theory. Suppose M is a model, $A = \mathrm{acl}(A)$, $B = \mathrm{acl}(B)$, and $A \cap B \subseteq M$. Let $C = \mathrm{acl}(AB)$ and let $(C_i)_{i<\omega}$ be an M-invariant Morley sequence over M with $C_0 = C$. Fix $\sigma \in \mathrm{Aut}(\mathbb{M}/M)$ with $\sigma(C_i) = C_{i+1}$ for all $i < \omega$. By restricting the sequence $(C_i)_{i<\omega}$ to a subtuple, we obtain an M-invariant Morley sequence $(B_i)_{i<\omega}$ with $B_0 = B$. Let $D = \mathrm{acl}((B_i)_{i<\omega})$. Let $\tilde{K} = (\mathrm{acl}((C_i)_{i<\omega}))_K$. Let $\{u_i : i < \alpha\}$ be a basis for M_V. Let $\{v_i : i < \beta\}$ complete this set to a basis for A_V and let $(w_{0,j})_{j<\gamma}$ complete it to a basis for $(B_0)_V$, then let $(w_{i,j})_{j<\gamma}$ be the set of vectors completing $\{u_i : i < \alpha\}$ to a basis for $(B_i)_V$ corresponding to the $(w_{0,j})_{j<\beta}$—i.e. $w_{i,j} = \sigma^i(w_{0,j})$. By our assumptions, $\{u_i : i < \alpha\} \cup \{v_i : i < \beta\} \cup \{w_{i,j} : i < \omega, j < \gamma\}$ is a set of linearly independent vectors in $\mathbb{M}_V$. Let $\tilde{V}$ be the $\tilde{K}$-vector space with basis $\{u_i : i < \alpha\} \cup \{v_i : i < \beta\} \cup \{w_{i,j} : i < \omega, j < \gamma\}$. To define the model $N = (\tilde{V}, \tilde{K})$, we are left with defining the form on $\tilde{V}$—for this it suffices to define the form on a basis. First, interpret the form so that N extends the structure on D—i.e.

$$[u_i, u_{i'}]^N = k \iff [u_i, u_{i'}]^D = k$$
$$[u_i, w_{i',j}]^N = k \iff [u_i, w_{i',j}]^D = k$$
$$[w_{i,j}, w_{i',j'}]^N = k \iff [w_{i,j}, w_{i',j'}]^D = k.$$

And likewise, interpret the structure so that it extends the structure on A—i.e.

$$[u_i, v_{i'}]^N = k \iff [u_i, v_{i'}]^A = k$$
$$[v_i, v_{i'}]^N = k \iff [v_i, v_{i'}]^A = k.$$

Then finally, we interpret the form so that the structure generated by AB_i does not depend on i: put $[v_i, w_{0,j}]^N = k \iff [v_i, w_{0,j}]^C = k$ and set

$$[v_i, w_{i',j}]^N = \begin{cases} k & \text{if } [v_i, w_{0,j}]^C = k \in A \\ \sigma^{i'}(k) & \text{if } [v_i, w_{0,j}]^C = k \notin A \end{cases}$$

This defines N. By quantifier-elimination, there is an embedding $\iota : N \to \mathbb{M}$ over D into $\mathbb{M}$. Let $A' = \iota(A)$. By quantifier-elimination, we have $AB_0 \equiv_M A'B_i$ for all i. This shows $\mathrm{tp}(A/B)$ does not Kim-divide over M. □

Proposition 3.9.36. *Suppose $M \models T_\infty$. Then*

1. *$a \mathop{\smile\!\!\!\!\!\!|}^{\Gamma}_{M} b \implies a \mathop{\smile\!\!\!\!\!\!|}^{K}_{M} b$.*

2. *There are a and b so that $a \mathop{\smile\!\!\!\!\!\!|}^{K}_{M} b$ and $a \not\mathop{\smile\!\!\!\!\!\!|}^{\Gamma}_{M} b$.*

Proof. (1) Suppose $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\Gamma}_{M} b$. By transitivity of $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\mathrm{ACF}}$, $(\mathrm{dcl}(aM))_K \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\mathrm{ACF}}_{M_K} (\mathrm{dcl}(bM))_K$ so

$$\mathrm{acl}(aM)_K \mathop{\smile\!\!\!\!\!\!\mid}\limits^{\mathrm{ACF}}_{M_K} (\mathrm{acl}(bM))_K$$

since the field points of the algebraic closure of any set X are just the field-theoretic algebraic closure of $(\mathrm{dcl}(X))_K$. Similarly, transitivity of independence for vector spaces forces $\langle (aM)_V \rangle \cap \langle (bM)_V \rangle \subseteq \langle M \rangle$. This shows $\mathrm{acl}(aM) \cap \mathrm{acl}(bM) = M$ so $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{K}_{M} b$, by Proposition 3.9.35.

(2) Given any $M \models T_\infty$, choose two vectors $b_1, b_2 \in \mathbb{M}_V$ that are $\mathbb{M}_K$-linearly independent over M. By model-completeness, we can find some vector a so that $\mathrm{acl}(aM) \cap \mathrm{acl}(b_1 b_2 M) \subseteq M$, so $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{K}_{M} b_1 b_2$, and also $[a, b_1] = [a, b_2]$. Then we clearly have $\{[a, b_1], [a, b_2]\}$ algebraically dependent, as they are equal, hence $a \not\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\Gamma}_{M} b_1 b_2$. $\square$

Remark 3.9.37. This observation implies that axioms (1)-(5) in Theorem 6.2.4 do not suffice to characterize $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{K}$, since $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\Gamma}$ satisfies these axioms and $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\Gamma} \neq \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{K}$ by Proposition 3.9.36(2).

Chapter 4

Local character

This chapter is joint work with Itay Kaplan and Saharon Shelah.

4.1 Introduction

A well-known theorem of Kim and Pillay characterizes the simple theories as those theories with an independence relation satisfying certain properties and shows that, moreover, any such independence relation must coincide with non-forking independence. As the theory of simple theories was being developed, work of Chatzidakis on ω-free PAC fields and Granger on vector spaces with bilinear forms furnished examples of non-simple theories for which there are nonetheless independence relations satisfying many of the fundamental properties of non-forking independence in simple theories. These properties include extension, symmetry, and the independence theorem. In Chapter 1, we proved an analogue of one direction of the Kim-Pillay theorem for NSOP_1 theories, showing essentially that the existence of an independence relation with these properties implies that a theory is NSOP_1. To establish the other direction, we introduced *Kim-independence* and showed that it is well-behaved in any NSOP_1 theory. The theory of Kim-independence provides an explanation for the simplicity-like phenomena observed in certain non-simple examples and a central issue of research concerning NSOP_1 theories is to determine the extent to which properties of non-forking independence in simple theories carry over to Kim-independence in NSOP_1 theories. This chapter addresses the specific issue of local character for Kim-independence.

Simple theories are defined to be the theories in which forking satisfies local character. Local character of non-forking asserts that there is some cardinal $\kappa(T)$ so that, for any complete type p over A, there is a set $B \subseteq A$ with $|B| < \kappa(T)$ over which p does not fork. An analogue of local character for Kim-independence in NSOP_1 theories was proved in Theorem III.4.5. It was shown there that if T is NSOP_1 and $M \models T$, then for any $p \in S(M)$, there is $N \prec M$ with $|N| < \kappa = \left(2^{|T|}\right)^+$ such that p does not Kim-fork over N.

However, this result was an unsatisfactory generalization of local character in simple theories for three reasons. First, with respect to non-forking, it follows almost immediately that if $\kappa(T)$ exists at all, it can be taken to be $|T|^+$: given a type $p \in S(A)$ with no $B \subseteq A$ of size $< |T|^+$ over which p does not fork, one can find a chain of forking types of length $|T|^+$ and then by the pigeonhole principle, some formula must fork infinitely often with respect to the same disjunction of dividing formulas. This equivalence is no longer immediate when considering Kim-independence, because of the added constraint that the formulas must divide with respect to Morley sequences and it was asked [KR17, Question 4.7] if $\left(2^{|T|}\right)^+$ can be replaced by $|T|^+$ in an arbitrary NSOP$_1$ theory. Secondly, non-forking independence satisfies *base monotonicity*, which means that if $p \in S(A)$ does not fork over B, then p does not fork over B' whenever $A \subseteq B' \subseteq B$. In other words, local character of forking implies that every type does not fork over an entire *cone* of small subsets of its domain. However, in an NSOP$_1$ theory T, Kim-independence satisfies base monotonicity if and only if T is simple. One would like an analogue of local character for NSOP$_1$ theories that shows that types over models do not Kim-divide over *many* small submodels. Finally, local character of non-forking independence *characterizes* simple theories. Many tameness properties of Kim-independence are known to characterize NSOP$_1$ theories, e.g. symmetry and the independence theorem, so it is natural to ask if local character does as well.

Our main theorem is:

Theorem 4.1.1. *Suppose T is a complete theory with monster model $\mathbb{M} \models T$. The following are equivalent:*

1. *T is NSOP$_1$.*
2. *There is no continuous increasing sequence of $|T|$-sized models $\left\langle M_i \,\middle|\, i < |T|^+ \right\rangle$ with union M and $p \in S(M)$ such that $p \restriction M_{i+1}$ Kim-forks over M_i for all $i < |T|^+$.*
3. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a stationary subset of $[M]^{|T|}$.*
4. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide contains a club subset of $[M]^{|T|}$.*
5. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a club subset of $[M]^{|T|}$.*
6. *Suppose that $N \models T$, $M \prec N$ and $p \in S(N)$ does not Kim-divide over M. Then the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a club subset of $[M]^{|T|}$.*

The equivalence of (1) and (2) was noted in Corollary III.4.6 with $|T|^+$ replaced by $\left(2^{|T|}\right)^+$, which is considerably weaker than the theorem proved here.

In particular, this theorem implies that if T is NSOP$_1$, $M \models T$, and $p \in S(M)$, then the set of $N \prec M$ with $|N| = |T|$ such that p does not Kim-fork over N is *non-empty.* However, by demanding a stronger form of local character, we obtain a new characterization of NSOP$_1$.

Remark 4.1.2. In the first draft of this chapter, published online as a preprint on July 2017, we did not yet have (5) or (6) above. Shortly after that draft was available, Pierre Simon have found an easier proof of (1) implies (4), and we thank him for allowing us to include his proof here. Later we found a proof of (6). These proofs uses symmetry of Kim-independence, but are not straightforward as in the proof in simple theories, and our original proof.

Our original proof assumes towards contradiction that local character fails and reaches a contradiction to NSOP_1 as is done in e.g. simple theories. For this approach to work we used stationary logic. This logic expands first-order logic by introducing a quantifier aa interpreted so that $M \models (\mathsf{aa}S)\,\varphi(S)$ if and only if the set of countable subsets $X \subseteq M$ such that M, when expanded with the predicate S interpreted as X, satisfies $\varphi(S)$ contains a club of $[M]^{\omega}$. This logic was introduced by Shelah in [She75] and later studied by Mekler and Shelah [MS86] who showed that the satisfiability of a theory in $L(\mathsf{aa})$ implies the satisfiability of a theory in a related logic, where the second-order quantifiers range over *uncountable* sets of a certain size. This theorem, which may be regarded as a version of the upward Lowenheim-Skolem theorem, provides a tool for "stretching" a family of counterexamples to local character in such a way that preserves the cardinality and continuity constraints needed to produce SOP_1.

After further review, we noticed that our original proof gives rise to a new phenomenon, which we call dual local character.

4.2 Preliminaries

NSOP_1 theories, invariant types, and Morley sequences

Definition 4.2.1. [DS04, Definition 2.2] A formula $\varphi(x;y)$ has the 1-*strong order property* (*SOP_1*) if there is a tree of tuples $\langle a_\eta \,|\, \eta \in 2^{<\omega}\rangle$ so that

- For all $\eta \in 2^{\omega}$, the partial type $\{\varphi(x;a_{\eta\restriction n}) \,|\, n < \omega\}$ is consistent.
- For all $\nu, \eta \in 2^{<\omega}$, if $\nu \frown \langle 0\rangle \trianglelefteq \eta$ then $\left\{\varphi(x;a_\eta), \varphi\left(x;a_{\nu\frown\langle 1\rangle}\right)\right\}$ is inconsistent.

A theory T is *$NSOP_1$* if no formula has SOP_1 modulo T.

Fact 4.2.2. Proposition III.2.4 T has NSOP_1 if and only if there is a formula $\varphi(x;y)$, $k < \omega$, and a sequence $\langle \bar{c}_i \,|\, i \in I\rangle$ with $\bar{c}_i = (c_{i,0}, c_{i,1})$ satisfying:

1. For all $i \in I$, $c_{i,0} \equiv_{\bar{c}_{<i}} c_{i,1}$.
2. $\{\varphi(x;c_{i,0}) \,|\, i \in I\}$ is consistent.
3. $\{\varphi(x;c_{i,1}) \,|\, i \in I\}$ is k-inconsistent.

We also use following notation. Write $a \mathop{\smile\!\!\!\!\!\vert}\,^{u}_{M} B$ for $\mathrm{tp}(a/MB)$ is finitely satisfiable in M, in other words it is a *coheir* of its restriction to M. A type $p \in S(M)$ is an *heir* of its restriction to $N \prec M$ if for every formula $\varphi(x;y) \in L(N)$ and every $b \in M$, if $\varphi(x;b) \in p$

then $\varphi(x;b') \in p$ for some $b' \in N$. We denote this by $c \mathop{\smile\!\!\!\!\!\!\mid}\,^{h}_{N} M$. This is equivalent to saying that $M \mathop{\smile\!\!\!\!\!\!\mid}\,^{u}_{N} c$.

Definition 4.2.3. A global type $q \in S(\mathbb{M})$ is called *A-invariant* if $b \equiv_A b'$ implies $\varphi(x;b) \in q$ if and only if $\varphi(x;b') \in q$. A global type q is *invariant* if there is some small set A such that q is A-invariant. If $q(x)$ and $r(y)$ are A-invariant global types, then the type $(q \otimes r)(x,y)$ is defined to be $\operatorname{tp}(a,b/\mathbb{M})$ for any $b \models r$ and $a \models q|_{\mathbb{M}b}$. It is also A-invariant. We define $q^{\otimes n}(x_0,\ldots,x_{n-1})$ by induction: $q^{\otimes 1} = q$ and $q^{\otimes n+1} = q(x_n) \otimes q^{\otimes n}(x_0,\ldots,x_{n-1})$.

Fact 4.2.4. [She90, Lemma 4.1] If T is any complete theory, $M \models T$, and $p \in S(M)$, then there is a complete global type q extending p which is, moreover, finitely satisfiable in M. In particular, q is M-invariant.

Definition 4.2.5. Suppose q is an A-invariant global type and I is a linearly ordered set. By a *Morley sequence in q over A of order type I*, we mean a sequence $\langle b_\alpha \mid \alpha \in I\rangle$ such that for each $\alpha \in I$, $b_\alpha \models q|_{Ab_{<\alpha}}$ where $b_{<\alpha} = \langle b_\beta \mid \beta < \alpha\rangle$. Given a linear order I, we will write $q^{\otimes I}$ for the A-invariant type in variables $\langle x_\alpha \mid \alpha < I\rangle$ so that for any $B \supseteq A$, if $\bar{b} \models q^{\otimes I}|_B$ then $b_\alpha \models q|_{Bb_{<\alpha}}$ for all $\alpha \in I$. If q is, moreover, finitely satisfiable in A, then we refer to a Morley sequence in q over A as a *coheir sequence* over A.

The above definition of $q^{\otimes I}$ generalizes the finite tensor product $q^{\otimes n}$ – given any global A-invariant type q and linearly ordered set I, one may easily show that $q^{\otimes I}$ exists and is A-invariant by compactness.

Definition 4.2.6. Suppose M is a model.

1. Given a formula $\varphi(x;b)$ and a global M-invariant type $q \supseteq \operatorname{tp}(b/M)$, we say that $\varphi(x;b)$ *k-Kim-divides over M via q* if, whenever $\langle b_i \mid i < \omega\rangle$ is a Morley sequence over M in q, then $\{\varphi(x;b_i) \mid i < \omega\}$ is k-inconsistent.

2. If q is a global M-invariant type with $q \supseteq \operatorname{tp}(b/M)$, we say $\varphi(x;b)$ *Kim-divides over M via q* if $\varphi(x;b)$ k-Kim-divides over M via q for some $k < \omega$.

3. We say $\varphi(x;b)$ Kim-divides over M if $\varphi(x;b)$ Kim-divides over M via q for some global M-invariant $q \supseteq \operatorname{tp}(b/M)$.

4. We say that $\varphi(x;b)$ Kim-forks over M if it implies a finite disjunction of formulas, each Kim-dividing over M.

5. We write $a \mathop{\smile\!\!\!\!\!\!\mid}\,^{K}_{M} B$ for $\operatorname{tp}(a/MB)$ does not Kim-fork (or Kim-independent) over M.

Note that if $a \mathop{\smile\!\!\!\!\!\!\mid}\,^{u}_{M} B$ then $a \mathop{\smile\!\!\!\!\!\!\mid}\,^{f}_{M} B$ (i.e. $\operatorname{tp}(a/BM)$ does not fork over M) which implies $a \mathop{\smile\!\!\!\!\!\!\mid}\,^{K}_{M} B$.

Fact 4.2.7. Theorem III.3.15 The following are equivalent for the complete theory T:

1. T is NSOP_1.

2. (Kim's lemma for Kim-dividing) Given any model $M \models T$ and formula $\varphi(x;b)$, $\varphi(x;b)$ Kim-divides via q for *some* global M-invariant $q \supseteq \operatorname{tp}(b/M)$ if and only if $\varphi(x;b)$ Kim-divides via q for *all* global M-invariant $q \supseteq \operatorname{tp}(b/M)$.

From this it easily follows that Kim-forking is equal to Kim-dividing Proposition III.3.19. The notion of Kim independence, denoted by $\perp^K$, satisfies many nice properties which turn out to be equivalent to NSOP$_1$.

Fact 4.2.8. Theorem III.8.1 The following are equivalent for the complete theory T:

1. T is NSOP$_1$.
2. Symmetry of Kim independence over models: $a \perp^K_M b$ iff $b \perp^K_M a$ for any $M \models T$.
3. Independence theorem over models: if $A \perp_M B$, $c \perp_M A$, $c' \perp_M B$ and $c \equiv_M c'$ then there is some $c'' \perp_M AB$ such that $c'' \equiv_{MA} c$ and $c'' \equiv c'_{MB}$.

Fact 4.2.9. Lemma III.7.6 Suppose that T is NSOP$_1$ and that $\langle a_i \mid i < \omega \rangle$ is an $\perp^K$-Morley sequence over M in the sense that $a_i \perp^K_M a_{<i}$ and the sequence is indiscernible. Then if $\varphi(x, a_0)$ does not Kim-divide over M, then $\{\varphi(x, a_i) \mid i < \omega\}$ does not Kim-divide over M, and in particular it is consistent.

The generalized club filter

Definition 4.2.10. Let κ be a cardinal and X a set with $|X| \geq \kappa$. We write $[X]^\kappa$ to denote $\{Y \subseteq X \mid |Y| = \kappa\}$.

1. A set $C \subseteq [X]^\kappa$ is *unbounded* if for every $Y \in [X]^\kappa$, there is some $Z \in C$ with $Y \subseteq Z$.
2. A set $C \subseteq [X]^\kappa$ is *closed* if, whenever $\langle Y_i \mid i < \alpha \leq \kappa \rangle$ is a chain in C, i.e. each $Y_i \in C$ and $i < j < \alpha$ implies $Y_i \subseteq Y_j$, then $\bigcup_{i<\alpha} Y_i \in C$.
3. A set $C \subseteq [X]^\kappa$ is *club* if it is closed and unbounded.
4. A set $S \subseteq [X]^\kappa$ is *stationary* if $S \cap C \neq \emptyset$ for every club $C \subseteq [X]^\kappa$.

The *club filter* on $[X]^\kappa$ is the filter generated by the clubs. If $|X| = \kappa$, then the club filter on $[X]^\kappa$ is the principal ultrafilter consisting of subsets of $[X]^\kappa$ containing X.

Example 4.2.11. If M is an L-structure of size $\geq \kappa \geq |L|$, then the collection of elementary substructures of M of size κ is a club in $[M]^\kappa$.

Remark 4.2.12. In the literature, e.g. [Jec13, Definition 8.21], the above definitions are given instead for subsets of $\mathcal{P}_{\kappa^+}(X) = \{Y \subseteq X \mid |Y| < \kappa^+\}$ but note that $[X]^\kappa$ is a club subset of $\mathcal{P}_{\kappa^+}(X)$, hence all definitions relativize to this set in the natural way.

Fact 4.2.13. Let κ be a cardinal and X a set with $|X| \geq \kappa^+$.

1. The club filter on $[X]^{\kappa}$ is κ^{+}-complete [Jec13, Theorem 8.22].

2. For every club $C \subseteq [X]^{\kappa}$, there is a collection of finitary functions $\overline{f} = \langle f_i : i < \kappa \rangle$ with $f_i : X^{n_i} \to X$ such that $C_{\overline{f}} := \{Y \in [X]^{\kappa} \mid f_i(Y^{n_i}) \subseteq Y \text{ for all } i < \kappa\}$ is contained in C. Equivalently, there is a function $F : X^{<\omega} \to [X]^{\kappa}$ such that $C_F \subseteq C$ [Jec13, Lemma 8.26].

3. Conversely, given a collection of finitary functions $\overline{f} = \langle f_i : i < \kappa \rangle$ with $f_i : X^{n_i} \to X$, the set $C_{\overline{f}}$ is club in $[X]^{\kappa}$.

4. When $\kappa = \omega$, for any club $C \subseteq [X]^{\kappa}$, there is a function $F : X^{<\omega} \to X$ such that $C_F \subseteq C$ [Jec13, Theorem 8.28].

We leave the proof of the next lemma to the reader.

Lemma 4.2.14. *Suppose λ is a cardinal, X is a set with $|X| = \lambda^{+}$, and $\langle Y_\alpha : \alpha < \lambda^{+} \rangle$ is an increasing continuous sequence of sets of cardinality λ with union X. Then $\{Y_\alpha \mid \alpha < \lambda^{+}\}$ is a club of $[X]^{\lambda}$. In particular, if $X = \lambda^{+}$ and $C \subseteq \lambda^{+} \setminus \lambda$ is a club of λ^{+}, then C is a club of $[X]^{\lambda}$.*

4.3 The main theorem

A short proof of (1) implies (4) in Theorem 4.1.1 using heirs

Here we give a short proof of (1) implies (4) in Theorem 4.1.1, due to Pierre Simon. We thank him for allowing us to include this proof.

Lemma 4.3.1. *Suppose $p(x) \in S(M)$, $M \models T$. Then the set of $N \prec M$ such that $|N| = |T|$ and p is an heir of $p|_N$ is a club subset of $[M]^{|T|}$.*

Proof. It is easy to verify that this set is closed under increasing unions, so it is enough to show that it contains a club.

Consider the L_p-structure M_p expanding M by forcing p to be definable — i.e. for every L-formula $\varphi(x; y)$ add a relation $R_\varphi(y)$ interpreted as $\left\{b \in M^{|y|} \,\middle|\, \varphi(x, b) \in p\right\}$. Note that $|L_p| = |L|$. Then if $N' \prec M_p$ then its L-reduct N is such that p is an heir of $p|_N$. Thus we are done by Example 4.2.11. □

Theorem 4.3.2. *Suppose T is NSOP$_1$. If $M \models T$ and $p \in S(M)$, then the set of elementary substructures $N \prec M$ with $|N| = |T|$ such that p does not Kim-divide over N contains a club.*

Proof. By Lemma 4.3.1, it suffices to show that if p is an heir of $p|_N$, then p does not Kim-divide over N. But if p is an heir of $p|_N$, then, given $c \models p$, $M \mathop{\smile\!\!\!\!\!|}^{u}_{N} c$, hence $M \mathop{\smile\!\!\!\!\!|}^{K}_{N} c$ by symmetry of Kim-independence (in fact one needs only a weak version of symmetry, see Proposition III.3.22) which implies $c \mathop{\smile\!\!\!\!\!|}^{K}_{N} M$. This shows that p does not Kim-divide over N. □

A proof of (1) implies (6) in Theorem 4.1.1

Lemma 4.3.3. *Suppose T is an arbitrary theory and $M \models T$ with $|M| \geq |T| = \kappa$. Given any global M-finitely satisfiable type q, let C_q denote the set of $N \prec M$ with $|N| = \kappa$ such that $q^{\otimes\omega}|_N = r^{\otimes\omega}|_N$ for some global N-finitely satisfiable type r. Then:*

1. *C_q is in the club filter on $[M]^\kappa$.*

2. *Given any set A, there is some $N \prec M$ of size $\leq |T| + |A|$ such that $A \subseteq N$ and $q^{\otimes\omega}|_N$ is a type of a Morley sequence generated by some global type r finitely satisfiable in N and if $\varphi(x, c)$ Kim-divides over M via q then $\varphi(x, c)$ Kim-divides over N via r.*

Proof. One proof of (1) essentially follows from the proof of Lemma III.4.4, so we also give an alternative one. Let $\bar{a} = \langle a_i : i < \omega \rangle$ be a coheir sequence generated by q over M. Then, $N \in C_q$ iff $N \prec M$ and $\bar{a}$ is a coheir sequence over N in the sense that $\mathrm{tp}(a_i / a_{<i} N)$ if finitely satisfiable in N. Thus it is easy to see that C_q is closed under unions.

Note that if $N \prec M$ is such that $\mathrm{tp}(\bar{a}/M)$ is an heir extension of its restriction to N, then $N \in C_q$: if $\varphi(a_i, a_{<i})$ holds when $\varphi(x, y)$ is some formula over N, then for some $c \in M$, $\varphi(c, a_{<i})$ holds, and by choice of N, we may assume that $c \in N$. Now Lemma 4.3.1 finishes the proof.

(2) is immediate from (1), applied to the theory $T(A)$ obtained from T by adding constants for the elements of A. □

Theorem 4.3.4. *Suppose T is NSOP$_1$ with $|T| = \kappa$ and $M \models T$. Then for a finite tuple b and any set A, the following are equivalent:*

1. *$A \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$.*

2. *There is a club $C \subseteq [M]^\kappa$ of elementary substructures of M such that $A \mathop{\smile\!\!\!\!\!\!\mid}^K_N b$ for all $N \in C$.*

3. *There is a stationary set $S \subseteq [M]^\kappa$ of elementary substructures of M such that $A \mathop{\smile\!\!\!\!\!\!\mid}^K_N b$ for all $N \in S$.*

Proof. (1) $\Longrightarrow$ (2) Suppose that $A \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$. Let $q \supseteq \mathrm{tp}(b/M)$ be a global M-finitely satisfiable type and choose $\langle b_i : i < \omega \rangle \models q^{\otimes\omega}|_M$ with $b_0 = b$. By Lemma 4.3.3, there is a club C_q of elementary substructures $N \prec M$ with $|N| = |T|$ so that $q^{\otimes\omega}|_N = r^{\otimes\omega}|_N$ for some global N-finitely satisfiable type r. Fix $N \in C$, a a finite tuple from A and $\varphi(x; b, n) \in \mathrm{tp}(a/Nb)$. As $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$, we know $\{\varphi(x; b_i, n) \mid i < \omega\}$ is consistent. As $\langle b_i : i < \omega \rangle$ is also a Morley sequence over N in a global N-finitely satisfiable type, it follows from Kim's lemma for Kim-dividing (Fact 4.2.7) that $\varphi(x; b, n)$ does not Kim-divide over N. As $\varphi(x; b, n)$ was arbitrary, we conclude $a \mathop{\smile\!\!\!\!\!\!\mid}^K_N b$. Since this was true for any a, we have that $A \mathop{\smile\!\!\!\!\!\!\mid}^K_N b$.

(2) $\Longrightarrow$ (3) is immediate.

(3) $\Longrightarrow$ (1) Suppose $a \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ for some finite tuple a from A. Let $\varphi(x; b, m) \in \mathrm{tp}(a/Mb)$ be a formula witnessing this. Fix $q \supseteq \mathrm{tp}(b/M)$ a global M-finitely satisfiable type and

$b_i : i < \omega\rangle \models q^{\otimes\omega}|_M$. Let $C' = \{N \prec M \mid |N| = |T| \text{ and } m \in N\}$. The set C' is clearly club so the intersection $C'' = C_q \cap C'$ is in the club filter on $[M]^\kappa$. If $N \in C''$ and $q^{\otimes\omega}|_N = r^{\otimes\omega}|_N$ for some global type r finitely satisfiable in N, then $\varphi(x; b, m) \in \mathrm{tp}(a/Nb)$ and $\langle b_i : i < \omega\rangle$ realizes $r^{\otimes\omega}|_N$. As $\{\varphi(x; b_i, m) \mid i < \omega\}$ is inconsistent, we have $a \not\mathrel{\raise0.2ex\hbox{$\smile$}\kern-0.7em\raise0.55ex\hbox{$|$}}^K_N b$. As S is stationary, it must intersect C'', so we get a contradiction. $\square$

Corollary 4.3.5. *Suppose T is NSOP$_1$ with $|T| = \kappa$ and $M \models T$. Then for a finite tuple a and any set B, the following are equivalent:*

1. *$a \mathop{\smile\!\!\!\!|}^K_M B$.*
2. *There is a club $C \subseteq [M]^\kappa$ of elementary substructures of M such that $a \mathop{\smile\!\!\!\!|}^K_N B$ for all $N \in C$.*
3. *There is a stationary set $S \subseteq [M]^\kappa$ of elementary substructures of M such that $a \mathop{\smile\!\!\!\!|}^K_N B$ for all $N \in S$.*

Proof. Follows immediately from symmetry of Kim-independence and Theorem 4.3.4. $\square$

Lemma 4.3.6. *Suppose T is NSOP$_1$. Assume $M \prec N$. Suppose that $a \mathop{\smile\!\!\!\!|}^K_M N$ and $\varphi(x, a)$ Kim-divides over N for $\varphi(x, y) \in L(M)$. Then $\varphi(x, a)$ Kim-divides over M.*

Proof. Let $\langle a_i : i < \omega\rangle$ be an indiscernible sequence over N starting with $a_0 = a$ such that $a_i \mathop{\smile\!\!\!\!|}^h_N a_{<i}$ and $\{\varphi(x, a_i) \mid i < \omega\}$ is inconsistent (to construct it, let $\langle b_i : i \in \mathbb{Z}\rangle$ be a coheir sequence in the type of $\mathrm{tp}(a/N)$, so in particular $b_i \mathop{\smile\!\!\!\!|}^u_N b_{<i}$ for $i < 0$, hence $b_{>i} \mathop{\smile\!\!\!\!|}^u_N b_i$ by transitivity of $\mathop{\smile\!\!\!\!|}^u$, and let $a_i = b_{-i}$ for $i < \omega$).

Then $\langle a_i \mid i < \omega\rangle$ is an $\mathop{\smile\!\!\!\!|}^K$-Morley sequence over M in the sense that $a_i \mathop{\smile\!\!\!\!|}^K_M a_{<i}$. To see this, suppose not, i.e., by symmetry suppose that $a_{<i} \not\mathop{\smile\!\!\!\!|}^K_M a_i$. Then for some formula $\psi(z, x)$ over M, $\psi(a_{<i}, a_i)$ holds and $\psi(z, a_i)$ Kim-divides over M. Since $a_{<i} \mathop{\smile\!\!\!\!|}^u_N a_i$, for some $n \in N$, $\psi(n, a_i)$ holds. However, since $a_i \equiv_N a$, by symmetry $N \mathop{\smile\!\!\!\!|}^K_M a_i$ — contradiction.

Suppose that $\varphi(x, a)$ does not Kim-divide over M. Then by Fact 4.2.9, $\{\varphi(x, a_i) \mid i < \omega\}$ is consistent — contradiction. $\square$

Lemma 4.3.7. *Suppose T is NSOP$_1$. Suppose that $\langle M_i \mid i \leq \alpha\rangle$ is an increasing sequence of elementary substructures of a model N, that $M_\alpha = \bigcup\{M_i \mid i < \alpha\}$ and that $p \in S(N)$. Assume that p does not Kim-fork over M_i for all $i < \alpha$. Then p does not Kim-fork over M_α.*

Proof. Let $a \models p$. We want to show that $a \mathop{\smile\!\!\!\!|}^K_{M_\alpha} N$, so by symmetry it is enough to show that $N \mathop{\smile\!\!\!\!|}^K_{M_\alpha} a$. Suppose not. Then there is some formula $\varphi(x, y)$ in $L(M_\alpha)$ and some $b \in N$ such that $\varphi(b, a)$ holds and $\varphi(x, a)$ Kim-divides over M_α. Let $i < \alpha$ be such that $\varphi(x, y) \in L(M_i)$. Since $M_\alpha \subseteq N$ and $a \mathop{\smile\!\!\!\!|}^K_{M_i} N$ by assumption, $a \mathop{\smile\!\!\!\!|}^K_{M_i} M_\alpha$. Hence by Lemma 4.3.6, $\varphi(x, a)$ Kim-divides over M_i. Hence $b \not\mathop{\smile\!\!\!\!|}^K_{M_i} a$. But this is a contradiction since $a \mathop{\smile\!\!\!\!|}^K_{M_i} N$ so by symmetry $b \mathop{\smile\!\!\!\!|}^K_{M_i} a$. $\square$

We can now prove (1) $\implies$ (6) from Theorem 4.1.1.

Theorem 4.3.8. *Suppose that T is NSOP$_1$. Suppose that a is a finite tuple, $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M N$ and $M \prec N$. Then the set E of $M' \in [M]^{|T|}$ such that $M' \prec M$ and $a \mathop{\smile\!\!\!\!\!\!\mid}^K_{M'} N$ is a club.*

Proof. The family E is closed under unions by Lemma 4.3.7. Hence to finish we only need to show that E contains a club, and this follows from Corollary 4.3.5 (1) $\implies$ (2). □

The equivalence (1)–(6)

We finish the proof of Theorem 4.1.1 with the following.

Theorem 4.3.9. *Suppose T is a complete theory. The following are equivalent:*

1. *T is NSOP$_1$.*
2. *There is no continuous increasing sequence of $|T|$-sized models $\langle M_i \,|\, i < |T|^+ \rangle$ with union M and $p \in S(M)$ such that $p \restriction M_{i+1}$ Kim-forks over M_i for all $i < |T|^+$.*
3. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a stationary subset of $[M]^{|T|}$.*
4. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide contains a club subset of $[M]^{|T|}$.*
5. *For any $M \models T$, $p \in S(M)$, the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a club subset of $[M]^{|T|}$.*
6. *Suppose that $N \models T$, $M \prec N$ and $p \in S(N)$ does not Kim-divide over M. Then the set of elementary substructures of M of size $|T|$ over which p does not Kim-divide is a club subset of $[M]^{|T|}$.*

Proof. (1) $\implies$ (6) is Theorem 4.3.8.

(6) $\implies$ (5) $\implies$ (4) $\implies$ (3) is trivial (for (6) implies (5), note that for $p \in S(M)$, p does not Kim-divide over M trivially).

(3) $\implies$ (2) By Lemma 4.2.14, $C = \{M_i \,|\, i < |T|^+\}$ is a club of $[M]^{|T|}$. As T is NSOP$_1$, there is a stationary set $S \subseteq [M]^{|T|}$ such that $N \in S$ implies p does not Kim-fork over N. Choose any $M_i \in C \cap S$ to obtain a contradiction.

(2) $\implies$ (1). Suppose T has SOP$_1$ as witnessed by some formula $\varphi(x, y)$. Let T^{sk} be a Skolemized expansion of T. Then T^{sk} also has SOP$_1$ as witnessed by $\varphi(x, y)$. Thus by Proposition 4.2.2, we can find a formula $\varphi(x, y)$ and an array $\langle c_{i,j} \,|\, i < \omega, j < 2 \rangle$ such that $c_{i,0} \equiv_{\bar{c}_{<i}} c_{i,1}$ for all $i < \omega$, $\{\varphi(x, c_{i,0}) \,|\, i < \omega\}$ is consistent and $\{\varphi(x; c_{i,1}) \,|\, i < \omega\}$ is 2-inconsistent (all in $\mathbb{M}^{sk}$). By Ramsey and compactness we may assume that $\langle \bar{c}_i \,|\, i < \omega \rangle$ is indiscernible (with respect to $\mathbb{M}^{sk}$) and extend this sequence to length $|T|^+$.

For $i \leq |T|^+$, let $N_i = \operatorname{dcl}(\bar{c}_{<i})$ (in $\mathbb{M}^{sk}$). Then for every limit ordinal $\delta < |T|^+$, $\varphi(x, c_{\delta,1})$ Kim-divides over N_δ as the sequence $\left\langle c_{j,1} \,\middle|\, \delta \leq j < |T|^+ \right\rangle$ is indiscernible and for all $\delta \leq j$, $\bar{c}_j \mathop{\smile\!\!\!\!\!\mid}\nolimits^u_{N_\delta} \bar{c}_{>j}$. As $c_{\delta,1} \equiv_{\bar{c}_{<\delta}} c_{\delta,0}$, it follows that $c_{\delta,1} \equiv_{N_\delta} c_{\delta,0}$, and hence $\varphi(x, c_{\delta,0})$ also Kim-divides. Let $p \in S\left(N_{|T|^+}\right)$ be any complete type containing $\{\varphi(x, c_{\delta,0}) \,|\, \delta < \kappa\}$, which is possible as this partial type is consistent. The sequence $\left\langle N_\delta \,\middle|\, \delta \in \lim\left(|T|^+\right)\right\rangle$ is an increasing and continuous sequence of elementary substructures of $N_{|T|^+}$ of size $|T|$ with union $N_{|T|^+}$ witnessing that (2) fails. ☐

Corollary 4.3.10. *Suppose T is NSOP$_1$, $M \models T$, $M \prec N$, and $p \in S(N)$. Then p does not Kim-fork over M iff for every κ with $|T| \leq \kappa \leq |M|$, the set of elementary substructures of M of size κ over which p does not Kim-divide is a club subset of $[M]^\kappa$.*

Proof. Suppose that p does not Kim-fork over M. Let $A \subseteq M$ be any subset of M of size κ and apply Theorem 4.1.1 to the theory $T(A)$ obtained from T by adding new constant symbols for the elements of A.

For the other direction, apply the left hand side with $\kappa = |T|$ and use Corollary 4.3.5. ☐

Corollary 4.3.11. *Suppose T is NSOP$_1$ and $M \models T$. Then given any set A, there is a club $E \subseteq [M]^{|T|+|A|}$ such that $N \in E$ iff $A \mathop{\smile\!\!\!\!\!\mid}\nolimits^K_N M$.*

Proof. Let $\kappa = |A| + |T|$. By Corollary 4.3.10, we know for each finite tuple a from A, there is a club $E_a \subseteq [M]^\kappa$ so that $N \in E_a$ iff $a \mathop{\smile\!\!\!\!\!\mid}\nolimits^K_N M$. Let $E = \bigcap_{a \in A} E_a$. As $|A| \leq \kappa$ and the club filter on $[M]^\kappa$ is κ^+-complete (Fact 4.2.13(1)), E is a club of $[M]^\kappa$. By the strong finite character of Kim-independence, we have $A \mathop{\smile\!\!\!\!\!\mid}\nolimits^K_N M$ iff $N \in E$. ☐

A sample application

Proposition 4.3.12. *Suppose T is NSOP$_1$ and $A \models T$. Given any set C, there is some $C' \supseteq C$ with $|C'| = |C| + |T|$ such that $C' \cap A$ is a model and $C' \mathop{\smile\!\!\!\!\!\mid}\nolimits^K_{A \cap C'} A$.*

Proof. Let $\kappa = |C| + |T|$. Let $C_0 = C$ and, by Corollary 4.3.11, we may let $E_0 \subseteq [A]^\kappa$ be a club of elementary substructures of A such that $N \in E_0$ implies $C_0 \mathop{\smile\!\!\!\!\!\mid}\nolimits^K_N A$. By induction, we will choose sets C_i, clubs $E_i \subseteq [A]^\kappa$, and models $X_i \prec A$ such that

1. $X_i \in \bigcap_{j \leq i} E_i$ and $C_i \cap A \subseteq X_i$.
2. $C_{i+1} = C_i \cup X_i$.
3. For all $N \in E_i$, we have $C_i \mathop{\smile\!\!\!\!\!\mid}\nolimits^K_N A$.

Given $\langle C_i, X_i, E_i \,|\, i \leq n \rangle$, let $C_{n+1} = C_n \cup X_n$. By Corollary 4.3.11, we may let $E_{n+1} \subseteq [A]^\kappa$ be a club such that $N \in E_{n+1}$ implies $C_{n+1} \mathop{\smile\!\!\!\!\!\mid}\nolimits^K_N A$. As

$$\{X \in [A]^\kappa \,|\, C_{n+1} \cap A \subseteq X\}$$

is a club of $[A]^\kappa$, we may choose $X_{n+1} \in \bigcap_{i \leq n+1} E_i$ containing $C_{n+1} \cap A$. This completes the induction.

Let $C_\omega = \bigcup_{i<\omega} C_i$. By construction, $C_\omega \cap A = \bigcup_{i<\omega} X_i$. As $i < j$ implies $X_i \subseteq X_j$, and $i \geq n$ implies $X_i \in E_n$, it follows that

$$C_\omega \cap A = \bigcup_{i \geq n} X_i \in E_n$$

for all n, as E_n is club. Also as each X_i is a model, this additionally shows that $C_\omega \cap A$ is a model. Moreover, if $c \in C_\omega$ is a finite tuple, there is some n so that $c \in C_n$, hence $c \mathop{\smile\hspace{-0.9em}\mid\hspace{0.4em}}^K_{C_\omega \cap A} A$, by the choice of E_n. Setting $C' = C_\omega$, we finish. □

4.4 A proof using stationary logic

More on clubs

Definition 4.4.1. Suppose κ is a cardinal and $A \subseteq B$, $S \subseteq [A]^\kappa$, and $T \subseteq [B]^\kappa$. We define $S^B \in [A]^\kappa$ and $T \upharpoonright A \in [A]^\kappa$ by

$$\begin{aligned} S^B &= \{Y \in [B]^\kappa \,|\, Y \cap A \in S\} \\ T \upharpoonright A &= \{X \in [A]^\kappa \,|\, \text{ there is } Y \in T \text{ such that } X = Y \cap A\}. \end{aligned}$$

Fact 4.4.2. [Jec13, Theorem 8.27] Suppose κ is a cardinal, $A \subseteq B$, $S \subseteq [A]^\kappa$, and $T \subseteq [B]^\kappa$.

1. If S is stationary in $[A]^\kappa$, then S^B is stationary in $[B]^\kappa$.
2. If T is stationary in $[B]^\kappa$, then $T \upharpoonright A$ is stationary in $[A]^\kappa$.

Lemma 4.4.3. *Suppose X is a set and λ and κ are cardinals with $\lambda \leq \kappa < |X|$. Suppose, moreover, we are given a stationary subset $\mathcal{S} \subseteq [X]^\kappa$ and, for every $Y \in \mathcal{S}$, a stationary subset $S_Y \in [Y]^\lambda$. Then $\mathcal{S}' = \bigcup_{Y \in \mathcal{S}} S_Y$ is a stationary subset of $[X]^\lambda$.*

Proof. Suppose $D \subseteq [X]^\lambda$ is a club. We must show $\mathcal{S}' \cap D \neq \emptyset$. By Fact 4.2.13(3), there is a sequence of finitary functions $\overline{f} = \langle f_i \,|\, i < \lambda\rangle$ where for all $i < \lambda$, $f_i : X^{n_i} \to X$ and the set $C_{\overline{f}} \subseteq [X]^\lambda$ of λ-sized subsets of X closed under $\overline{f}$ is a club with $C_{\overline{f}} \subseteq D$. The subsets of X of size κ closed under $\overline{f}$ form a club $C^*_{\overline{f}} \subseteq [X]^\kappa$, hence $C^*_{\overline{f}} \cap \mathcal{S} \neq \emptyset$. Fix $Y \in C^*_{\overline{f}} \cap \mathcal{S}$. Define a sequence of functions $\overline{g} = \langle g_i \,|\, i < \lambda\rangle$ by $g_i = f_i \upharpoonright Y^{n_i}$ for all $i < \lambda$. This definition makes sense as Y is closed under the functions f_i so that $C_{\overline{f}} \cap [Y]^\lambda = C_{\overline{g}}$, the subsets of Y closed under $\overline{g}$, hence is a club of $[Y]^\lambda$. Therefore $C_{\overline{f}} \cap [Y]^\lambda \cap S_Y \neq \emptyset$. In particular, this shows $D \cap S' \neq \emptyset$, which completes the proof. □

The club filter on $[X]^\omega$ was characterized by Kueker in terms of games of length ω [Kue72]. The natural analogue for games of length λ determines a filter on $\mathcal{P}_{\lambda^+}(X)$, which, in general, differs from the club filter. In generalizing stationary logic to quantification over sets of some uncountable size λ, it turns out that this filter provides a more useful analogue to the club filter on $[X]^{\omega}$ than the club filter on $[X]^{\lambda}$.

Definition 4.4.4. Suppose X is a set and λ is a regular cardinal. Given a subset $F \subseteq \mathcal{P}_{\lambda^+}(X)$, we define the game $G(F)$, to be the game of length λ where Players I and II alternate playing an increasing λ sequence of elements of $\mathcal{P}_{\lambda^+}(X)$. In this game, Player II wins if and only if the union of the sets played is in F. The filter $D_\lambda(X)$ is defined to be the filter generated by the sets $F \subseteq \mathcal{P}_{\lambda^+}(X)$ in which Player II has a winning strategy in $G(F)$. We say $Y \subseteq \mathcal{P}_{\lambda^+}(X)$ is $D_\lambda(X)$-*stationary* if Y intersects every set in $D_\lambda(X)$.

It is easy to check that every club $C \subseteq [X]^{\lambda}$ is an element of $D_\lambda(X)$ and, therefore, that every $S \subseteq [X]^{\lambda}$ that is $D_\lambda(X)$-stationary is also stationary with respect to the usual club filter on $[X]^{\lambda}$. It was remarked in [MS86] that if $\lambda = \lambda^{<\lambda}$, then $D_\lambda(\lambda^+)$ is just the filter generated by the clubs of λ^+ intersected with the set of ordinals of cofinality λ (considered as initial segments of λ^+). More precisely, we have the following fact. (We omit its proof since it is not necessary for the rest.)

Fact 4.4.5. Suppose λ is an infinite cardinal and write $S^{\lambda^+}_{\lambda}$ for the stationary set $\{\alpha < \lambda^+ \mid \mathrm{cf}(\alpha) = \lambda\}$.

1. If $C \subseteq \lambda^+$ is a club, then $C \cap S^{\lambda^+}_{\lambda} \in D_\lambda(\lambda^+)$.
2. Suppose $\lambda = \lambda^{<\lambda}$. Then $D_\lambda(\lambda^+)$ is generated by sets of the form $C \cap S^{\lambda^+}_{\lambda}$, where $C \subseteq \lambda^+$ is a club.

Lemma 4.4.6. *Suppose X is a set of size λ^+, and $\langle X_\alpha \mid \alpha < \lambda^+\rangle$ is an increasing and continuous sequence from $\mathcal{P}_{\lambda^+}(X)$ with union X. Suppose $S \subseteq \mathcal{P}_{\lambda^+}(X)$ is $D_\lambda(X)$-stationary. Then the set $S_* = \{\alpha < \lambda^+ \mid cf(\alpha) = \lambda, X_\alpha \in S\}$ is a stationary subset of λ^+.*

Proof. As $|X| = \lambda^+$, we may assume $X = \lambda^+$. Let $C \subseteq \lambda^+$ consist of the ordinals $\alpha < \lambda^+$ such that $X_\alpha = \alpha$. This set is easily seen to be a club.

Let $C_* \subseteq \lambda^+$ be a club. We must show $C_* \cap S_* \neq \emptyset$. By Fact 4.4.5(1), $C_* \cap C \cap S^{\lambda^+}_{\lambda} \in D_\lambda(X)$, hence $\left(S^{\lambda^+}_{\lambda} \cap C \cap C_*\right) \cap S \neq \emptyset$. Pick Y in this intersection. Then by definition of C, $Y = X_\alpha = \alpha$ for some $\alpha \in S^{\lambda^+}_{\lambda}$. As $X_\alpha \in S$, we have $\alpha \in S_*$. This shows $S_* \cap C_* \neq \emptyset$. □

Lemma 4.4.7. *Suppose $A \subseteq B$ and $S \subseteq \mathcal{P}_{\lambda^+}(B)$ is $D_\lambda(B)$-stationary. Then the set $S \restriction A = \{X \cap A \mid X \in S\}$ is $D_\lambda(A)$-stationary.*

Proof. It is enough to show that if $F \in D_\lambda(A)$ then $F^B = \{X \in \mathcal{P}_{\lambda^+}(B) \mid X \cap A \in F\} \in D_\lambda(B)$. We may assume that there is some winning strategy f for Player II in the game $G(F)$, since $F \in D_\lambda(A)$. That is, the function f is defined so that if, at stage i, Player I has played $\langle A_j \mid j \leq i\rangle$ then $f(\langle A_j \mid j \leq i\rangle)$ outputs the play for Player II.

Now we will define a winning strategy for Player II in the game $G\left(F^B\right)$. At stage i, if Player I has played $\langle A_j \mid j \leq i\rangle$, Player II plays $B_i = A_i \cup f\left(\langle A_j \cap A \mid j \leq i\rangle\right)$. As the rules of the game require that the sets are increasing, we have

$$A_i \cap A \subseteq f\left(\langle A_j \cap A \mid j \leq i\rangle\right) \subseteq A,$$

hence $B_i \cap A = f\left(\langle A_j \cap A \mid j \leq i\rangle\right)$. It follows that

$$\begin{array}{c|ccccc} \mathrm{I} & A_0 \cap A & & A_1 \cap A & & \cdots \\ \hline \mathrm{II} & & B_0 \cap A & & B_1 \cap A & \cdots \end{array}$$

is a play according to f in $G(F)$. Therefore,

$$\left(\bigcup_{i<\lambda} A_i \cup B_i\right) \cap A = \bigcup_{i<\lambda} (A_i \cap A) \cup (B_i \cap A) \in F,$$

which shows $\bigcup_{i<\lambda} A_i \cup B_i \in F^B$. We have shown that Player II has a winning strategy in $G\left(F^B\right)$ so $F^B \in D_\lambda(B)$. $\square$

Stationary logic

The stationary logic $L(\mathsf{aa})$ was introduced in [She75] (where it was called $L\left(Q^{ss}_{\aleph_1}\right)$). The logic is defined as follows: given a first-order language L, expand the language with countably many new unary predicates $\{S_i \mid i < \omega\}$ and a new quantifier aa. The formulas of L in $L(\mathsf{aa})$ are the the smallest class containing the first-order formulas of L, closed under the usual first-order formation rules together with the rule that if φ is a formula, then $(\mathsf{aa}S_i)\,\varphi$ is also a formula, for any new unary predicate S_i. Satisfaction is defined as usual, together with the rule that $M \models (\mathsf{aa}S)\,\varphi(S)$ if and only if $M \models \varphi(S)$ when $S^M = X$ for "almost all" $X \in [M]^\omega$—that is, $\left\{X \in [M]^\omega \,\middle|\, \text{if } S^M = X \text{ then } M \models \varphi(S)\right\}$ contains a club of $[M]^\omega$. We define the quantifier stat dually: $M \models (\mathsf{stat}S)\,\varphi(S)$ if and only if $M \models \neg(\mathsf{aa}S)\,\neg\varphi(S)$. Note that $M \models (\mathsf{stat}S)\,\varphi(S)$ if and only if $\left\{X \in [M]^\omega \,\middle|\, \text{if } S^M = X \text{ then } M \models \varphi(S)\right\}$ is stationary. Given an L-structure M, we write $\mathrm{Th}_{\mathsf{aa}}(M)$ for the set of $L(\mathrm{aa})$-sentences satisfied by M. We refer the reader to [BKM78, Section 1] for a detailed treatment of stationary logic.

Later work by Mekler and Shelah extended stationary logic, which quantifies over *countable sets*, to a logic that permits quantification over sets of higher cardinality [MS86]. For λ a regular cardinal, the logic $L\left(\mathsf{aa}^\lambda\right)$ is defined analogously to $L(\mathsf{aa})$, with semantics defined so that $M \models \left(\mathsf{aa}^\lambda S\right)\varphi(S)$ if and only if $\left\{X \in [M]^\lambda \,\middle|\, \text{if } S^M = X \text{ then } M \models \varphi(S)\right\} \in D_\lambda(M)$. The quantifier stat^λ is also understood dually: $M \models \left(\mathsf{stat}^\lambda S\right)\varphi(S)$ if and only if $M \models \neg\left(\mathsf{aa}^\lambda S\right)\neg\varphi(S)$. If T is an $L(\mathsf{aa})$-theory, one obtains an $L\left(\mathsf{aa}^\lambda\right)$-theory by replacing the quantifier aa with aa^λ. We call this theory the *λ-interpretation of T*. By working with $D_\lambda(M)$ instead of the full club filter on $[M]^\lambda$, one is able to relate satisfiability of an $L(\mathsf{aa})$-theory to the satisfiability of its λ-interpretation. Below, the "moreover" clause about λ-saturation is not stated in [MS86], but is immediate from the proof.

Fact 4.4.8. [MS86, Theorem 1.3] Suppose $\lambda = \lambda^{<\lambda}$ and T is a consistent $L(\mathsf{aa})$-theory of size at most λ. Then the λ-interpretation of T has a model of size at most λ^+. In fact, there is such a model which is, moreover, λ-saturated.

The following easy observation is also useful:

Lemma 4.4.9. *Suppose φ is a first-order formula, possibly with parameters from M and $|\varphi(M)| > \aleph_0$. Then if $M' \models Th_{\mathsf{aa}}(M)$ in the λ-interpretation, then $|\varphi(M')| > \lambda$.*

Proof. Suppose not. Then $\left\{S \in [M']^\lambda \,\middle|\, \varphi(M') \subseteq S\right\}$ is a club of $[M']^\lambda$ hence an element of $D_\lambda(M')$. Therefore $M' \models \left(\mathsf{aa}^\lambda S\right) \forall x (\varphi(x) \to S(x))$. As $M' \models \mathrm{Th}_{\mathsf{aa}}(M')$ in the λ-interpretation, $M \models (\mathsf{aa}S) \forall x (\varphi(x) \to S(x))$, so $\varphi(M)$ is countable, a contradiction. □

Reduction to a countable language

Remark 4.4.10. Suppose that T is an NSOP$_1$ theory in the language L. Suppose that $M \models T$ and $\varphi(x, y)$ is any formula. Then for any language $L' \subseteq L$ containing φ, and any $b \in \mathbb{M}$, $\varphi(x, b)$ Kim-divides over M in L iff $\varphi(x, b)$ Kim-divides over $M' := M \restriction L'$ (in the sense of $T \restriction L'$). Indeed, this follows from Kim's lemma for Kim-dividing (Fact 4.2.7) and the fact that if $\bar{b}$ is a coheir sequence in L over M starting with b, then it is also in L'.

Lemma 4.4.11. *Suppose T is an NSOP$_1$ theory in the language L, $M \models T$ and for some $p \in S(M)$, the set*

$$S = \{N \prec M \mid |N| = |T|,\ p \text{ Kim-divides over } N\}$$

is stationary in $[M]^{|T|}$. Then there is a countable sublanguage $L' \subseteq L$ and a stationary set $S' \subseteq [M]^\omega$ so that, setting $p' = p \restriction L$, we have that for all $N' \in S'$, p' Kim-divides over N'.

Proof. For each $N \in S$, choose some $\varphi_N(x; b_N) \in p$ such that $\varphi(x; b_N)$ Kim-divides over N. By Fact 4.2.13(1), the club filter on $[M]^{|T|}$ is $|T|^+$-complete, so for any a partition of a stationary set into $|T|$ many pieces, we may find some piece which is stationary. Therefore we may assume there is some formula φ so that $\varphi_N(x; b_N) = \varphi(x; b_N)$ for all $N \in S$. Let L' be any countable sublanguage of L containing φ. By Remark 4.4.10 and (the proof of) Theorem 4.3.4, for each $N \in S$, there is a club $C_N \subseteq [N]^\omega$ of countable L'-elementary substructures over which $\varphi(x; b_N)$ Kim-divides. By Lemma 4.4.3, $S' = \bigcup_{N \in S} C_N$ is a stationary subset of $[M]^\omega$. By definition of S', if $N' \in S'$, then there is some $\varphi(x; b_{N'}) \in p$ such that $\varphi(x; b_{N'})$ Kim-divides over N'. □

Stretching

Lemma 4.4.12. *Suppose T is NSOP$_1$, $|T| = \aleph_0$, $M \models T$, and there is $p \in S(M)$ so that the set*

$$S_0 = \{N \prec M \mid |N| = \aleph_0 \text{ and } p \text{ Kim-divides over } N\}$$

is stationary. Then, given any regular uncountable cardinal $\lambda = \lambda^{<\lambda}$*, there is a model* $M' \models T$*,* $|M'| = \lambda^+$*, a formula* $\varphi(x; y)$*, and a type* p_* *over* M' *so that*

$$S'_0 = \{N' \prec M' \mid |N'| = \lambda, \text{ there is } \varphi(x; a'_N) \in p_* \text{ that Kim-divides over } N'\}$$

is $D_\lambda(M')$*-stationary.*

Proof. As no type Kim-divides over its domain, it follows that M is uncountable. For each $N \in S_0$, there is some formula $\varphi_N(x; a_N) \in p$ and $k_N < \omega$ so that $\varphi(x; a_N)$ k_N-Kim-divides over N via a Morley sequence in some global N-finitely satisfiable type. As the club filter on $[M]^\omega$ is $\aleph_1$-complete, Fact 4.2.13(1), there are φ and k so that for some stationary $S \subseteq S_0$, we have $N \in S$ implies $\varphi_N(x; a_N) = \varphi(x; a_N)$ and $k_N = k$.

Let $l = |a_N|$ for all $N \in S$ and let $\tilde{M}$ be an $\aleph_1$-saturated elementary extension of M. Let χ be a sufficiently large regular cardinal so that all objects of interest are contained in $H(\chi)$. In particular, we may choose χ so that $\tilde{M}, {}^\omega\tilde{M}, T$, L, and p are contained in $H(\chi)$, together with a bijection to ω witnessing the countable cardinality of L, and we consider the structure

$$\mathcal{H} = \left(H(\chi), \in, M, \tilde{M}, L, T, p\right).$$

By Fact 4.4.2(1), the set $S_* = \{X \in [\mathcal{H}]^\omega \mid X \cap M \in S\}$ is a stationary subset of $[\mathcal{H}]^\omega$.

Let $\Phi(X)$ be the formula in the language of $\mathcal{H}$ together with a new predicate X that naturally asserts: there exists $c \in M^l$, such that $\varphi(x; c) \in p$ and such that there exists $f \in {}^\omega\left(\tilde{M}^l\right))$ such that:

$X \cap M$ is an elementary substructure of M.

$f = \langle f_i \mid i < \omega\rangle$ is an $(X \cap M)$-indiscernible sequence such that $\operatorname{tp}(f_i / (X \cap M) f_{<i})$ is finitely satisfiable in $(X \cap M)$.

$f(0) = c$.

$\{\varphi(x; f_i) \mid i < \omega\}$ is k-inconsistent.

We first show the following:

Claim. $\mathcal{H} \models (\mathsf{stat}X)\,\Phi(X)$.

Proof of claim. As S_* is stationary, it suffices to show that if $X \in S_*$ and $S^{\mathcal{H}} = S_*$ then $\mathcal{H} \models \Phi(S)$. Recall that if $X \in S_*$, then $X \cap M \in S$ so $X \cap M$ is a countable elementary substructure of M, and $\varphi(x; a_{X\cap M})$ is a formula in p that k-Kim-divides over $X \cap M$. As $\tilde{M}$ is $\aleph_1$-saturated, there is a coheir sequence $\langle a_i \mid i < \omega\rangle$ over $X \cap M$ in $\tilde{M}$ with $a_0 = a_{X\cap M}$ and $\{\varphi(x; a_i) \mid i < \omega\}$ k-inconsistent. Put $c = a_0$ and let $f \in {}^\omega(\tilde{M}^l)$ be defined by $f_i = a_i$, we easily have (1)-(4) satisfied, proving the claim. □

By Fact 4.4.8, there is $\mathcal{H}'$ which is a model of the λ-interpretation of $\mathrm{Th}_{aa}(\mathcal{H})$ with $|\mathcal{H}'| = \lambda^+$, $\mathcal{H}' = \left(\mathcal{H}', \in', M', \tilde{M}', L', p'\right)$. As L and T are coded by natural numbers, the language L is contained in L' and thus the definable set $\left\{x \in \mathcal{H}' \,\middle|\, x \in \tilde{M}'\right\}$ may be regarded as the domain of an L'-structure whose reduct to L is a model of T and likewise for M. Moreover

$M' \prec_L \tilde{M}'$ and $|M'| = \lambda^+$, by Lemma 4.4.9. As $\mathcal{H}' \models (\mathsf{stat}^\lambda X)\, \Phi(X)$, there is a $D_\lambda(\mathcal{H}')$-stationary set $S'_* \subseteq [\mathcal{H}']^\lambda$ witnessing this. Let $S' = S'_* \upharpoonright M'$—i.e. $S' = \{X \cap M' \mid X \in S'_*\}$. By Lemma 4.4.7, S' is $D_\lambda(M')$-stationary. Let $p_* = p' \upharpoonright L$. To conclude the proof, it suffices to establish the following:

Claim. p_* is a type over M' and if $N \in S'$, then p_* k-Kim-divides over N via some $\varphi(x; a'_N) \in p_*$.

Proof of claim. It is clear that p_* is a consistent type over M'. Now fix $N \in S'$. By definition of S', $N = X \cap M'$ for some $X \in [\mathcal{H}']^\lambda$ such that $\mathcal{H}' \models \Phi(S)$ when $S^{\mathcal{H}'} = X$. It follows that for some $b \in M'$, there is an N-indiscernible sequence $\langle b_i \mid i \in I\rangle$ with $b_0 = b$, such that $\operatorname{tp}(b_i/Nb_{<i})$ is finitely satisfiable in N, $\varphi(x; b_0) \in p'$ and $\{\varphi(x; b_i) \mid i \in I\}$ is k-inconsistent, where I denotes the (possibly non-standard) natural numbers of $\mathcal{H}'$. By indiscernibility, $\langle b_i \mid i \in I\rangle$ is a Morley sequence over N in a global N-finitely satisfiable type, which shows $\varphi(x; b_0)$ k-Kim-divides over N. This completes the proof. □

The main lemma

Lemma 4.4.13. *(Main Lemma) Suppose T is a complete theory, $M \models T$ is a model with $|M| \geq |T|$, and for some $p \in S(M)$, the set*

$$S_0 = \left\{ N \in [M]^{|T|} \,\middle|\, N \prec M, p \text{ Kim-divides over } N \right\}$$

is stationary. Then T has SOP_1.

Proof. Towards contradiction suppose T is NSOP_1. By Lemma 4.4.11, there is a countable sublanguage $L' \subseteq L$ and a stationary set $S'_0 \subseteq [M]^\omega$ such that if $p' = p \upharpoonright L'$ then for all $N \in S$, $N \prec_{L'} M$ and p' Kim-divides over N. Therefore, we may assume for the rest of the proof that T is countable.

By forcing with the LÃľvy collapse $\mathrm{Coll}(\lambda^+, 2^\lambda)$ for a sufficiently large cardinal, we may assume there is some an uncountable cardinal $\kappa = \kappa^{<\kappa}$, namely $\kappa = \lambda^+$, while preserving the situation. By Lemma 4.4.12, there is a model $M' \models T$ with $|M'| = \kappa^+$ and a type p' over M' so that

$$S''_0 = \left\{ N \in [M']^\kappa \,\middle|\, N \prec M' \text{ and some } \varphi(x; c_N) \in p' \text{ Kim-divides over } N \right\}$$

is $D_\kappa(M')$-stationary. Let $\langle M_\alpha \mid \alpha < \kappa^+\rangle$ be a continuous and increasing sequence of κ-sized elementary substructures of M' with union M'. The set $S = \{\alpha < \kappa^+ \mid \mathrm{cf}(\alpha) = \kappa, M_\alpha \in S''_0\}$ is a stationary subset of κ^+ by Lemma 4.4.6. By intersecting with a club, we may also assume that for all $\alpha \in S$, M_α contains c_{M_β} for all $\beta \in \alpha \cap S$.

From here, the proof closely follows the proof of Theorem III.4.5. For each $\alpha \in S$, let c_α denote c_{M_α} and let r_α be a global M_α-finitely satisfiable type extending $\operatorname{tp}(c_\alpha/M_\alpha)$. By reducing S, we may assume that there is some $k < \omega$ such that r witnesses that $\varphi(x; c_\alpha)$ k-Kim-divides over M_α. For each $\alpha \in S$, apply Lemma 4.3.3(2) to choose a countable $N_\alpha \prec M_\alpha$ such that $r_\alpha^{\otimes\omega}|_{N_\alpha}$ is the type of a Morley sequence in some global N_α-finitely satisfiable type and hence such that $\varphi(x; c_\alpha)$ k-Kim-divides over N_α. Define $\rho : S \to \kappa^+$ by $\rho(\alpha) = \min\{\beta < \alpha \,|\, N_\alpha \subseteq M_\beta\}$. This is well-defined and pressing down on S as κ is regular and uncountable. By Fodor's lemma, there is $S' \subseteq S$ such that ρ is constant on S', say with constant value β_0. As $|M_{\beta_0}| = \kappa$, there are $\leq \kappa^{\aleph_0} = \kappa$ many choices for $N_\alpha \subseteq M_{\beta_0}$ so there is a stationary $S'' \subseteq S'$ and N'_0 so that $N_\alpha = N'_0$ for all $\alpha \in S''$. As there are $\leq 2^{\aleph_0} \leq \kappa$ choices for $r_\alpha^{\otimes\omega}|_{N'_0}$, there is a stationary $S_* \subseteq S''$ such that $r_\alpha^{\otimes\omega}|_{N'_0}$ is constant, with value $s_0^{\otimes\omega}|_{N'_0}$ for some global coheir s_0 over N'_0. Let $\delta_0 = \min S_0$, $e_0 = c_{\delta_0}$.

Repeating this process ω many times, we find an increasing sequence $\langle \delta_i \,|\, i < \omega\rangle$ of ordinals in κ^+, an increasing sequence of models $\langle N'_i \,|\, i < \omega\rangle$, $e_i \in M'$ for $i < \omega$ and global N'_i-finitely satisfiable types s_i such that:

N'_i contains $e_{<i}$, $\varphi(x; e_j)$ is k-Kim-dividing over N'_i for every $i \leq j$, s_i is a global coheir over N'_i extending $\operatorname{tp}(e_i/N'_i)$ and $e_j \equiv_{N'_i} e_i$ for all $j \geq i$. In addition, $s_j^{\otimes\omega}|_{N'_i} = s_i^{\otimes\omega}|_{N'_i}$ for all $j \geq i$.

Denote $\bar{e} = \langle e_i : i < \omega\rangle$. Note that $\{\varphi(x; e_i) \,|\, i < \omega\}$ is a subset of p', hence consistent. Now, exactly as in the claim in the proof of Theorem III.4.5, we can show that if $i_0 < \ldots < i_{n-1} < \omega$ and for each $j < n$, $f_j \models s_{i_j}|_{N'_{i_j}\bar{e}f_{>j}}$ then $e_{i_j} \equiv_{e_{i_{<j}}f_{<j}} f_j$ for all $j < n$ and $\{\varphi(x; f_j) \,|\, j < n\}$ is k-inconsistent. By compactness, we can find an array $\langle (c_{i,0}, c_{i,1}) \,|\, i < \omega\rangle$ so that $\{\varphi(x, c_{i,0}) \,|\, i < \omega\}$ is consistent, $\{\varphi(x, c_{i,1}) \,|\, i < \omega\}$ is k-inconsistent, and $c_{i,0} \equiv_{\bar{c}_{<i}} c_{i,1}$ for all $i < \omega$. By Fact 4.2.2, we obtain SOP_1, a contradiction. $\square$

Corollary 4.4.14. *Theorem 4.1.1 (1) $\implies$ (4) holds.*

4.5 Dual local character

Definition 4.5.1. (T any theory) Say that a formula $\varphi(x, a)$ *strongly Kim-divides* over a model M if for every global M-invariant type $q \supseteq \operatorname{tp}(a/M)$, $\varphi(x, a)$ Kim-divides *over M via q*.

Remark 4.5.2. By Fact 4.2.7, strong Kim-dividing = Kim-dividing iff T is NSOP_1.

Definition 4.5.3. A *dual type* (over A) in x is a set F of (A-)definable sets in x such that for some $k < \omega$, it is k-inconsistent. Say that F *dually divides* over a model N, if every $X \in F$ which is not definable over N divides over N. Similarly define when F *dually Kim-divides* over N and when F *strongly dually Kim-divides* over N.

Theorem 4.5.4. *The following are equivalent for a complete theory T.*

1. *T is $NSOP_1$.*

2. *There is no continuous increasing sequence of $|T|$-sized models $\left\langle M_i \,\middle|\, i < |T|^+ \right\rangle$ with union M and a dual type F over M such that $F \restriction M_{i+1}$ does not strongly dually Kim-divide over M_i for all $i < |T|^+$.*

3. *Assume that $M \models T$ and F a dual type over M. Then there is a stationary subset S of $[M]^{|T|}$ such that if $N \in S$ then $N \prec M$ and F strongly dually Kim-divides over N.*

4. *(Dual local character) Same as (3) but S is a club.*

Proof. The proof is essentially dualizing or inverting the proof (using stationary logic) of Theorem 4.1.1 (1) $\implies$ (4), but we go into some details.

(1) $\implies$ (4). We follow the proof of "(1) implies (4)" of Theorem 4.1.1 as described in Section 4.4. Namely, assume that (2) fails. This means that there is a stationary subset S of $[M]^{|T|}$ such that if $N \in S$ then $N \prec M$ and there is some $X \in F$ which is not definable over N but still does not Kim-divide over N. Using the same proof as in Lemma 4.4.11, we may assume that the language L is countable and that there is a single formula $\varphi(x, y)$ with $|x| = n$ such that if $N \in S$ then for some $b \in M \backslash N$, $\varphi(x, b)$ does not Kim-divide over N (and $\varphi(x, b)$ is not N-definable). Now we repeat the same procedure as in Lemma 4.4.12. Thus, for a regular uncountable cardinal $\lambda = \lambda^{<\lambda}$, we get a model $M' \models T$, $|M'| = \lambda^+$, a formula $\varphi(x, y)$, and a k-inconsistent family F_* of definable formulas over M' so that

$$S'_0 = \{N' \prec M' \,|\, |N'| = \lambda, \exists \varphi(x; a'_N) \in F_* \text{ not } N'\text{-definable and does not Kim-divide over } N'\}$$

is $D_\lambda(M')$-stationary. Now we repeat the proof of Lemma 4.4.13. The contradiction we will arrive at the end will be the same contradiction, but the roles of the sequences e_i and f_j are reversed. Now $\{\varphi(x, e_i) \,|\, i < \omega\}$ is k-inconsistent (note that the formulas $\varphi(x, e_i)$ must define distinct definable sets from F_*) and $\langle \varphi(x, f_j) \,|\, j < n \rangle$ is consistent.

(4) $\implies$ (3) $\implies$ (2) is exactly as in the proof of Lemma 4.3.9. The proof of (2) $\implies$ (1) is just dualizing the proof of "(2) implies (1)" in Theorem 4.3.9 in the sense that the sequences $\langle c_{i,0} \,|\, i < \omega \rangle$ and $\langle c_{i,1} \,|\, i < \omega \rangle$ exchange places. $\square$

Question 4.5.5. *Is there a proof of the dual local character which does not use stationary logic? Such a proof may reveal some new properties of Kim-dividing.*

Chapter 5

SOP_1 in one variable

5.1 Introduction

This chapter is concerned with showing strong order property 1 (SOP_1) is witnessed by a formula with only one free variable. SOP_1 was introduced by Džamonja and Shelah in their study of the $\trianglelefteq^*$-order and they observed the class of NSOP_1 contains the simple theories [DS04]. Our subsequent work with Artem Chernikov in Chapter 1characterized NSOP_1 in terms of independent amalgamation of types, which gave a Kim-Pillay-style criterion for NSOP_1 that, in turn, implied that many nonsimple examples of interest lie within this class. Later in Chapter 3, with Itay Kaplan, we introduced the theory of Kim-independence which provided evidence that NSOP_1 is a meaningful dividing line, admitting a structure theory close to simplicity theory.

SOP_1 is distinctive among dividing lines because of the difficulty of showing that a theory is NSOP_1 directly by syntactic means. In essentially all known examples of non-simple NSOP_1 theories, one first shows that the theory has a well-behaved notion of independence and then makes use of the Kim-Pillay-style criterion from Corollary I.4.1 to show that this implies the theory is NSOP_1. In algebraic examples, such as Frobenius fields or bilinear forms over an algebraically closed field, this strategy is natural and closely parallels the established strategy for showing the simplicity of similar theories, such as bounded PAC fields or ACFA. However, in combinatorial examples, this approach can seem rather cumbersome or indirect.

We simplify the syntax of SOP_1 by proving that SOP_1 is always witnessed by a formula in a single free variable. One-variable theorems have been proved for almost all of the major dividing lines, both because it makes it easier to check whether a theory has the given property, and because it is a natural test question for one's understanding of the dividing line's behavior. Yet these theorems can sometimes be difficult to discover. For example, the questions of whether there are one-variable theorems for the strict order property or the tree property of the second kind were both posed as open problems by Shelah; the former was settled later by Lachlan [Lac75], the latter much later by Chernikov [Che14]. In some cases, the analysis can be simplfied by considering generalized indiscernibles, e.g. indiscernible arrays for TP_2 or indiscernible trees for $\mathrm{TP}_1/\mathrm{SOP}_2$, but these are of little direct use in studying formulas witnessing SOP_1 (see [HS14, p. 29n1] for a discussion). The argument below instead makes use of an equivalent formulation of SOP_1 in terms of a sequence of pairs to conclude by a direct combinatorial argument.

5.2 The proof

We begin by noting some equivalent formulations of SOP_1 in terms of arrays that will be useful. In referring to an array $(c_{i,j})_{i<\omega,j<2}$, we write $\bar{c}_i = (c_{i,0}, c_{i,1})$ and $\bar{c}_{<i} = (\bar{c}_k)_{k<i}$. We write $L(C)$ to denote the collection of L-formulas with parameters from the set C. We always assume T is a complete theory with monster model $\mathbb{M} \models T$.

Fact 5.2.1. Proposition III.2.4 The following are equivalent:

1. T has SOP_1—that is, there is a formula $\varphi(x; y)$ and a tree of tuples $(a_\eta)_{\eta\in 2^{<\omega}}$ so that
   a) For all $\eta \in 2^\omega$, $\{\varphi(x; a_{\eta|\alpha}) : \alpha < \omega\}$ is consistent.
   b) For all $\nu \frown \langle 0\rangle \trianglelefteq \eta \in 2^{<\omega}$, $\{\varphi(x; a_\eta), \varphi(x; a_{\nu\frown\langle 1\rangle})\}$ is inconsistent.
2. There is a formula $\varphi(x; y)$ and array $(c_{i,j})_{i<\omega,j<2}$ so that
   a) $c_{i,0} \equiv_{\bar{c}_{<i}} c_{i,1}$ for all $i < \omega$.
   b) $\varphi(x; c_{i,0}) : i < \omega\}$ is consistent.
   c) $\{\varphi(x; c_{i,1}) : i < \omega\}$ is 2-inconsistent.
3. There is an array $(c_{i,j})_{i<\omega,j<2}$, with $c_{i,j} = (d_{ij}, e_{ij})$ for all i, j, and formulas $\chi_1(x; y)$ and $\chi_2(x; z)$ so that, writing $\psi(x; y, z)$ for $\chi_1(x; y) \wedge \chi_2(x; z)$, the following conditions are satisfied:
   a) For all $i < \omega$, $e_{i,0} \equiv_{Cc_{<i,0}e_{<i,1}} e_{i,1}$;
   b) $\{\psi(x; c_{i,0}) : i < \omega\}$ is consistent;
   c) $j \leq i \implies \{\chi_1(x; d_{i,0}), \chi_2(x; e_{j,1})\}$ is inconsistent.

Remark 5.2.2. Although conditions (1)-(3) are not, in general, equivalent at the level of formulas, if one of the conditions is true for a formula $\varphi(x; y)$, then for any of the other conditions, there is a formula $\varphi'(x'; y')$ witnessing this with $l(x) = l(x')$. Hence we say T has SOP_1 witnessed by a formula in a single free variable if there is a $\varphi(x; y)$ with $l(x) = 1$ for at least one of the conditions (1)-(3).

Lemma 5.2.3. *If T has SOP_1, there is a formula $\varphi(x; y)$ and an array $(c_{i,0}, c_{i,1})_{i<\omega}$ so that*

1. *$\{\varphi(x; c_{i,0}) : i < \omega\}$ is consistent.*
2. *$\{\varphi(x; c_{i,1}) : i < \omega\}$ is 2-inconsistent.*

3. *$(\overline{c}_i)_{i<\omega}$ is an indiscernible sequence.*

4. *$c_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}$ for all $i < \omega$.*

5. *$(c_{k,0})_{k\geq i}$ is $\overline{c}_{<i}c_{i,1}$-indiscernible.*

Proof. If T has SOP_1, then by Fact 5.2.1(2), Ramsey, and compactness, there is a formula $\varphi(x;y)$ and an array $(b_{i,0}, b_{i,1})_{i<\omega}$ satisfying (1)-(4) in the statement. Define $c_{i,0} = b_{2i+1,0}$ and $c_{i,1} = b_{2i,1}$ for all $i < \omega$. The array $(c_{i,0}, c_{i,1})_{i<\omega}$ clearly satisfies (1)-(3). (5) is also clear, since $(b_{k,0})_{k\geq 2i+1}$ is $\overline{b}_{<2k+1}$-indiscernible. To see (4), note $b_{2i+1,0} \equiv_{\overline{b}_{<2i}} b_{2i+1,1}$ and hence $b_{2i+1,1} \equiv_{\overline{b}_{<2i}} b_{2i,1}$ by (3) so, by definition, $c_{i,0} \equiv_{\overline{c}_{<i}} c_{i,1}$. □

Lemma 5.2.4. *Suppose T does not witness SOP_1 with any formula in the variables x. Suppose b is a tuple of the same length as x, C is some set of parameters, and $(c_{i,0}, c_{i,1})_{i<\omega}$ is an array satisfying*

1. *$(\overline{c}_i)_{i<\omega}$ is a C-indiscernible sequence.*

2. *$(c_{i,0})_{i<\omega}$ is Cb-indiscernible.*

3. *$c_{i,0} \equiv_{C\overline{c}_{<i}} c_{i,1}$ for all $i < \omega$.*

4. *$(c_{k,0})_{k\geq i}$ is $C\overline{c}_{<i}c_{i,1}$-indiscernible.*

Then

$$tp(b/C(c_{i,0})_{i<\omega}) \cup \{\varphi(x;c_{i,0}) \leftrightarrow \varphi(x;c_{i,1}) : i < \omega, \varphi \in L(C)\}$$

is consistent.

Proof. Suppose not. Let N be maximal so that

$$\mathrm{tp}(b/C(c_{i,0})_{i<\omega}) \cup \{\varphi(x;c_{i,0}) \leftrightarrow \varphi(x;c_{i,1}) : i < N, \varphi \in L(C)\}$$

is consistent. By compactness, we may fix $\chi(x;c_{\leq M,0}) \in \mathrm{tp}(b/C(c_{i,0})_{i<\omega})$, a finite $\Delta(x) \subseteq \{\varphi(x;c_{i,0}) \leftrightarrow \varphi(x;c_{i,1}) : i < N, \varphi \in L(C)\}$, and a formula $\varphi \in L(C)$ so that

$$\chi(x;c_{\leq M,0}) \wedge \bigwedge \Delta(x) \vdash \varphi(x;c_{N,0}) \leftrightarrow \neg\varphi(x;c_{N,1}).$$

Necessarily by the choice of N, we have $M \geq N$, and, without loss of generality, $\chi(x;c_{\leq M,0}) \vdash \varphi(x;c_{N,0})$. Put $C' = C \cup \overline{c}_{<N}$ and let

$$\begin{aligned} a_{i.0} &= (c_{(M-N)i+N,0}, c_{(M-N)i+N+1,0}, \ldots, c_{(M-N)i+M,0}) \\ b_{i,0} &= c_{(M-N)i+N,0} \\ a_{i,1} &= c_{(M-N)i+N,1}. \end{aligned}$$

Unravelling definitions, we have $b_{i,0} \equiv_{C'a_{<i,0}\bar{b}_{<i}} b_{i,1}$ for all i. Therefore, we may choose, for all $i < \omega$, some $a_{i,1}$ so that $a_{i,0}b_{i,0} \equiv_{C'a_{<i,0}\bar{b}_{<0}} a_{i,1}b_{i,1}$. Set $z = (z_0, \ldots, z_{M-N})$ and define $\psi(x; z) \in L(C')$ by

$$\psi(x; z) = \chi(x; c_{<N,0}, z_0, \ldots, z_{M-N}) \wedge \bigwedge \Delta(x).$$

Write $\mu(x; y, z) = \varphi(x; y) \wedge \psi(x; z) \in L(C')$. Setting $d_{i,j} = (a_{i,j}, b_{i,j})$, for $i < \omega, j < 2$, we obtain an array $(d_{i,0}, d_{i,1})_{i<\omega}$ so that $d_{i,0} \equiv_{C'd_{<i,0}b_{<i,1}} d_{i,1}$.

Since $\chi(x; c_{<N,0}, c_{N,0}, \ldots, c_{M,0}) \in \text{tp}(b'/C(c_{i,0})_{i<\omega})$ and $(c_{i,0})_{i<\omega}$ is Cb-indiscernible, we have

$$\{\chi(x; c_{<N,0}, c_{(M-N)i+N,0}, \ldots, c_{(M-N)i+M,0}) : i < \omega\} \subseteq \text{tp}(b/C(c_{i,0})_{i<\omega}).$$

By construction, $\Delta(x) \cup \text{tp}(b/C(c_{i,0})_{i<\omega})$ is consistent. Unravelling definitions, we have $\{\mu(x; d_{i,0}) : i < \omega\}$ is consistent. By our choice of N, we know

$$\chi(x; c_{<N,0}, c_{N,0}, \ldots, c_{M,0}) \wedge \bigwedge \Delta(x) \wedge \varphi(x; c_{N,0}) \wedge \varphi(x; c_{N,1})$$

is inconsistent. By indiscernibility of the sequence $(\bar{c}_i)_{i<\omega}$, for all $i < \omega$,

$$\chi(x; c_{<N,0}, c_{(M-N)i+N,0}, \ldots, c_{(N-M)i+M,0}) \wedge \bigwedge \Delta(x) \wedge \varphi(x; c_{(M-N)i+N,0}) \wedge \varphi(x; c_{(M-N)i+N,1})$$

is inconsistent. Then as the sequence $(c_{k,0})_{k\geq(M-N)i+N}$ is $C\bar{c}_{<(M-N)i+N}c_{(M-N)i+N,1}$-indiscernible, it follows that, for all $j \geq i$,

$$\chi(x; c_{<N,0}, c_{(M-N)j+N,0}, \ldots, c_{(N-M)j+M,0}) \wedge \bigwedge \Delta(x) \wedge \varphi(x; c_{(M-N)j+N,0}) \wedge \varphi(x; c_{(M-N)i+N,1})$$

is inconsistent. Unravelling definitions again, this shows that for $i \leq j$, $\mu(x; d_{j,0}) \wedge \varphi(x; b_{i,1})$ is inconsistent. By Fact 5.2.1(3), this shows T has SOP$_1$ (in some formula in the variables x). $\square$

Theorem 5.2.5. *If T has SOP$_1$, there is some formula in a single free-variable with SOP$_1$.*

Proof. Suppose $\varphi(x, y; z)$ is a formula witnessing SOP$_1$. So there is an array $(c_{i,0}, c_{i,1})_{i<\omega}$ satisfying conditions (1)-(5) of Lemma 5.2.3 with respect to $\varphi(x, y; z)$. We will suppose T does not witness SOP$_1$ in the free variables y and we will exhibit a formula witnessing SOP$_1$ in the free variables x. Let C_n be the set enumerated by $\bar{c}_{<n}$. For $n \leq m < \omega$, define the partial type $q_{n,m}(y)$ by

$$q_{n,m}(y) = \{\psi(y; c_{m,0}) \leftrightarrow \psi(y; c_{m,1}) : \psi \in L(C_n)\}.$$

By induction on $n < \omega$, we will choose b_n so that

1. $\{\varphi(x, b_n; c_{i,0}) : i < \omega\}$ is consistent.

2. b_n realizes $q_{k,m}$ for all $k < n$ and $m \geq k$.

3. $(\overline{c}_k)_{k\geq n}$ is $C_n b_n$-indiscernible.

To begin, choose an arbitrary $(a_0, b_0) \models \{\varphi(x, y; c_{i,0}) : i < \omega\}$. By Ramsey, compactness, and automorphism, we may assume $(\overline{c}_i)_{i<\omega}$ is b_0-indiscernible. Next, suppose we are given b_n so that $\{\varphi(x, b_n; c_{i,0}) : i < \omega\}$ is consistent, b_n realizes $q_{k,m}$ for all $k < n$ and $m \geq k$, and $(\overline{c}_k)_{k\geq n}$ is $b_n C_n$-indiscernible. We additionally know $c_{k,0} \equiv_{C_n \overline{c}_{<k}} c_{k,1}$ and $(c_{l,0})_{l\geq k}$ is $C_n \overline{c}_{<k} c_{k,1}$-indiscernible for all $k \geq n$. Therefore, we may apply Lemma 5.2.4 to conclude

$$\mathrm{tp}(b_n/C_n(c_{i,0})_{i<\omega}) \cup \{\psi(y; c_{m,0}) \leftrightarrow \psi(y; c_{m,1}) : m \geq n, \psi \in L(C_n)\}$$

is consistent. Let b_{n+1} realize this partial type. By Ramsey, compactness, and an automorphism over C_n, we may assume $(\overline{c}_k)_{k\geq n}$ is $C_n b_{n+1}$-indiscernible, hence $(\overline{c}_k)_{k\geq n+1}$ is $C_{n+1} b_{n+1}$-indiscernible. It follows, then, that $\{\varphi(x, b_{n+1}; c_{i,0}) : i < \omega\}$ is consistent, $(\overline{c}_k)_{k\geq n+1}$ is $b_{n+1} C_{n+1}$-indiscernible, and b_{n+1} realizes $q_{k,m}(y)$ for all $k < n+1$ and $m \geq k$.

By compactness, then, we obtain a tuple b so that $\{\varphi(x, b; c_{i,0}) : i < \omega\}$ is consistent and b realizes $q_{n,k}(y)$ for all $n \leq k < \omega$. It follows that $c_{n,0} \equiv_{\overline{c}_{<n} b} c_{n,1}$ for all n. Setting $d_{i,j} = (b, c_{i,j})$ for all $i < \omega, j < 2$, we obtain an array $(d_{i,0}, d_{i,1})_{i<\omega}$ so that $\{\varphi(x; d_{i,0}) : i < \omega\}$ is consistent, $\{\varphi(x; d_{i,1}) : i < \omega\}$ is inconsistent, and $d_{i,0} \equiv_{\overline{d}_{<i}} d_{i,1}$ for all i. This shows there is a formula in the variables x witnessing SOP_1. This shows we may reduce the number of variables in the formula witnessing SOP_1 and, by induction, we conclude. $\square$

Part IV

Examples and Applications

Chapter 6

Generic expansion and Skolemization

This chapter is joint work with Alex Kruckman.

6.1 Introduction

Many of the early developments in the study of simple theories were guided by the thesis that a simple theory can be understood as a stable theory plus some 'random noise.' This loose intuition became a concrete recipe for creating new simple theories: start with a stable theory, and, through some kind of generic construction, add additional random structure in an expanded language. This strategy was pursued by Chatzidakis and Pillay [CP98], who showed that adding a generic predicate or a generic automorphism to a stable theory results in a simple theory which is, in general, unstable. In the case of adding a generic predicate, it suffices to assume that the base theory is simple; that is, expansion by a generic predicate preserves simplicity. The paper [CP98] spawned a substantial literature on generic structures and simple theories, which in turn shed considerable light on what a general simple theory might look like.

We are interested in using generic constructions to produce new examples of NSOP_1 theories. The class of NSOP_1 theories, which contains the class of simple theories, was isolated by Džamonja and Shelah [DS04] and later investigated by Shelah and Usvyatsov [SU08]. Until recently, very few non-simple examples were known to lie within this class. A criterion, modeled after the well-known theorem of Kim and Pillay characterizing the simple theories as those possessing a well-behaved independence relation, was observed in Chapter 1. This criterion was applied to show that the theory of an ω-free PAC field of characteristic zero and the theory of an infinite dimensional vector space over an algebraically closed field with a generic bilinear form are both NSOP_1. That chapter also showed, by a variation on a construction of Baudisch, that a simple theory obtained as a Fraïssé limit with no algebraicity may be 'parametrized' to produce an NSOP_1 theory which is, in general, non-simple. Chapter 3 developed a general theory of independence in NSOP_1 theories, called *Kim-independence*, which turns out to satisfy many of the familiar properties of forking independence in simple theories (e.g. extension, symmetry, the independence theorem, etc.). In this chapter, we apply this theory of independence to verify that certain generic constructions preserve NSOP_1.

In Section 6.2, we review the theory of Kim-independence in NSOP_1 theories and make some technical contributions to this theory. We establish strengthenings of the extension property, the chain condition, and the independence theorem for Kim-independence, obtaining additional instances of algebraic independence in their conclusions (see Definition 6.2.8, and Theorems 6.2.15, 6.2.18, and 6.2.21). The main deficiency of Kim-independence is the failure of base monotonicity, and this work can be viewed as an effort to circumvent that deficiency, since the instances of algebraic independence that we need would be automatic in the presence of base monotonicity (see Remarks 6.2.9 and 6.2.10).

Section 6.3 is dedicated to an analysis of the theory $T^{\emptyset}_{L}$ of the generic L-structure (the model completion of the empty theory in an arbitrary language L). The work in this section was motived by a preprint of Jeřábek [Jeř17]. In an early draft of [Jeř17], Jeřábek showed that $T^{\emptyset}_{L}$ is always NSOP_3, regardless of the language. He asked if this could be improved to NSOP_1 and if $T^{\emptyset}_{L}$ weakly eliminates imaginaries. We give positive answers to these questions, and we characterize Kim-independence and forking independence in this theory. In a subsequent draft of [Jeř17], Jeřábek also independently answered both questions.

But Jeřábek's first question suggested a much more general one. An L-theory T may be considered as an L'-theory for any language L' that contains L. A theorem of Winkler [Win75] establishes that, as an L'-theory, the theory T has a model completion $T_{L'}$, provided that T is model complete and eliminates the quantifier $\exists^{\infty}$. The theory $T_{L'}$ axiomatizes the generic expansion of T by the new constants, functions, and relations of L'. Using the theory developed in Section 6.2, we show that if T is NSOP_1, then $T_{L'}$ is as well; that is, generic expansions preserve NSOP_1.

In [Win75], Winkler also showed that if T is a model complete theory eliminating the quantifier $\exists^{\infty}$, then T has a generic Skolemization T_{Sk}. More precisely, if T is an L-theory, one may expand the language by adding a function f_{φ} for each formula $\varphi(\overline{x}, y)$ of L. And T, together with axioms asserting that each $f_{\varphi}(\overline{x})$ acts as a Skolem function for $\varphi(\overline{x}, y)$, has a model companion. This result was used by Nübling in [Nüb04], who showed that one may Skolemize *algebraic* formulas in a simple theory while preserving simplicity. Nübling further observed that, in general, adding a generic Skolem function for a non-algebraic formula produces an instance of the tree property, and hence results in a non-simple theory. We show, however, that generic Skolemization preserves NSOP_1. By iterating, we show that any NSOP_1 theory eliminating the quantifier $\exists^{\infty}$ can be expanded to an NSOP_1 theory with built-in Skolem functions, and we also characterize Kim-independence in the expansion in terms of Kim-independence in the original theory. This result is of intrinsic interest, but it also provides a new technical tool in the study of Kim-independence in NSOP_1 theories, which, at least at its current stage of development, only makes sense when the base is a model. Preservation of NSOP_1 by generic expansion and generic Skolemization is established in Section 6.4.

6.2 NSOP$_1$ and independence

Preliminaries on NSOP$_1$

Throughout this section, we fix a complete theory T and a monster model $\mathbb{M} \models T$.

Definition 6.2.1. A formula $\varphi(x; y)$ has SOP 1 modulo T if there is a tree of tuples $(a_\eta)_{\eta \in 2^{<\omega}}$ in $\mathbb{M}$ so that:

- For all $\eta \in 2^\omega$, the partial type $\{\varphi(x; a_{\eta|\alpha}) \mid \alpha < \omega\}$ is consistent.
- For all $\nu, \eta \in 2^{<\omega}$, if $\nu \frown \langle 0 \rangle \trianglelefteq \eta$ then $\{\varphi(x; a_\eta), \varphi(x; a_{\nu \frown \langle 1 \rangle})\}$ is inconsistent.

The theory T is NSOP$_1$ if no formula has SOP 1 modulo T. An incomplete theory is said to be NSOP$_1$ if every completion is NSOP$_1$.

Definition 6.2.2. We call any $p \in S(\mathbb{M})$ a *global type*. A global type p is *A-invariant* if, for all formulas $\varphi(x; y)$, if $b \equiv_A b'$, then $\varphi(x; b) \in p$ if and only if $\varphi(x; b') \in p$ (equivalently, p is invariant under the action of $\mathrm{Aut}(\mathbb{M}/A)$ on $S(\mathbb{M})$). If p is a global A-invariant type, a *Morley sequence in p over A* is a sequence $(b_i)_{i \in I}$ from $\mathbb{M}$ so that $b_i \models p|_{A(b_j)_{j<i}}$. We denote by $p^{\otimes n}|_A$ the type $\mathrm{tp}(b_{i_1}, \ldots, b_{i_n}/A)$ when $i_1 < i_2 < \cdots < i_n$. By invariance, this type does not depend on the choice of Morley sequence $(b_i)_{i \in I}$ or indices i_k.

Definition 6.2.3. Fix a model $M \prec \mathbb{M}$.

1. A formula $\varphi(x; b)$ *Kim-divides* over M if there is an M-invariant global type $q \supseteq \mathrm{tp}(b/M)$ so that if $(b_i)_{i<\omega}$ is a Morley sequence over M in q, then $\{\varphi(x; b_i) \mid i < \omega\}$ is inconsistent.
2. A partial type $p(x)$ *Kim-divides* over M if $p(x)$ implies some formula which Kim-divides over M.
3. We write $a \mathop{\smile\!\!\!\!\!\!|}\,^K_M b$ for the assertion that $\mathrm{tp}(a/Mb)$ does not Kim-divide over M.

A well-known theorem of Kim and Pillay characterizes the simple theories as those theories with a notion of independence satisfying certain properties—this serves both as a useful way to establish that a theory is simple and as a method to characterize forking independence for the given theory. An analogous criterion for establishing that a theory is NSOP$_1$ was proved in Chapter 1. Later, it was observed in Chapter 3 that this criterion gives rise to an abstract characterization of $\mathop{\smile\!\!\!\!\!\!|}\,^K$.

Theorem 6.2.4. *Proposition I.5.8, Theorem III.9.1 Assume there is an* $\mathrm{Aut}(\mathbb{M})$*-invariant ternary relation* $\mathop{\smile\!\!\!\!\!\!|}$ *on small subsets of* $\mathbb{M}$ *satisfying the following properties, for an arbitrary* $M \prec \mathbb{M}$ *and arbitrary tuples from* $\mathbb{M}$*:*

1. *Strong finite character: if $a \not\mathop{\smile\!\!\!\!\!\!|}_M b$, then there is a formula $\varphi(x, b, m) \in tp(a/Mb)$ such that for any $a' \models \varphi(x, b, m)$, $a' \not\mathop{\smile\!\!\!\!\!\!|}_M b$.*

2. *Existence over models:* $a \mathop{\smile\!\!\!\!\!\!|}_M M$.

3. *Monotonicity: if* $aa' \mathop{\smile\!\!\!\!\!\!|}_M bb'$*, then* $a \mathop{\smile\!\!\!\!\!\!|}_M b$.

4. *Symmetry: if* $a \mathop{\smile\!\!\!\!\!\!|}_M b$*, then* $b \mathop{\smile\!\!\!\!\!\!|}_M a$.

5. *The independence theorem: if* $a \mathop{\smile\!\!\!\!\!\!|}_M b$*,* $a' \mathop{\smile\!\!\!\!\!\!|}_M c$*,* $b \mathop{\smile\!\!\!\!\!\!|}_M c$ *and* $a \equiv_M a'$*, then there exists* a'' *with* $a'' \equiv_{Mb} a$*,* $a'' \equiv_{Mc} a'$*, and* $a'' \mathop{\smile\!\!\!\!\!\!|}_M bc$.

Then T *is NSOP*$_1$ *and* $\mathop{\smile\!\!\!\!\!\!|}$ *strengthens* $\mathop{\smile\!\!\!\!\!\!|}^K$*, i.e. if* $a \mathop{\smile\!\!\!\!\!\!|}_M b$*, then* $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$*. If* $\mathop{\smile\!\!\!\!\!\!|}$ *satisfies the following additional property, then* $\mathop{\smile\!\!\!\!\!\!|} = \mathop{\smile\!\!\!\!\!\!|}^K$ *over models, i.e.* $a \mathop{\smile\!\!\!\!\!\!|}_M b$ *if and only if* $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$*:*

6. *Witnessing: if* $a \not\mathop{\smile\!\!\!\!\!\!|}_M b$*, then there exists a formula* $\varphi(x, b, m) \in \mathrm{tp}(a/Mb)$*, such that for any Morley sequence* $(b_i)_{i<\omega}$ *over* M *in a global* M*-finitely satisfiable type extending* $tp(b/M)$*,* $\{\varphi(x, b_i, m) \mid i < \omega\}$ *is inconsistent.*

The following is a reformulation of some of the main results of Chapter 3:

Theorem 6.2.5. *If* T *is* NSOP$_1$*, then* $\mathop{\smile\!\!\!\!\!\!|}^K$ *satisfies the properties (1)-(6) in Theorem 6.2.4, as well as*

7. *Extension: if* $a \mathop{\smile\!\!\!\!\!\!|}_M b$*, then for any* c*, there exists* a' *such that* $a' \equiv_{Mb} a$ *and* $a' \mathop{\smile\!\!\!\!\!\!|}_M bc$.

8. *The chain condition: if* $a \mathop{\smile\!\!\!\!\!\!|}_M b$ *and* $I = (b_i)_{i<\omega}$ *is a Morley sequence over* M *in a global* M*-invariant type extending* $\mathrm{tp}(b/M)$*, then there exists* a' *such that* $a' \equiv_{Mb} a$*,* $a' \mathop{\smile\!\!\!\!\!\!|}_M I$*, and* I *is* Ma'*-indiscernible.*

We will also be interested in the relation of algebraic independence, $\mathop{\smile\!\!\!\!\!\!|}^a$. Algebraic independence comes close to satisfying the criteria in Lemma 6.2.4 in any theory, but it typically does not satisfy the independence theorem.

Definition 6.2.6. For any set $C \subset \mathbb{M}$ and any tuples a and b, we define

$$a \mathop{\smile\!\!\!\!\!\!|}^a_C b \iff \mathrm{acl}(Ca) \cap \mathrm{acl}(Cb) = \mathrm{acl}(C).$$

Lemma 6.2.7. *In any theory* T*,* $\mathop{\smile\!\!\!\!\!\!|}^a$ *satisfies extension, existence over models, monotonicity, symmetry, strong finite character, and witnessing.*

Proof. Extension for algebraic independence is proved in [Hod93, Theorem 6.4.5]. Existence over models, monotonicity, and symmetry are immediate from the definitions.

For strong finite character and witnessing, note that if $M \models T$, $a \not\mathop{\smile\!\!\!\!\!\!|}^a_M b$, and $c \in (\mathrm{acl}(Ma) \cap \mathrm{acl}(Mb)) \setminus M$ witnesses this, then one can choose $\chi(z; a, m) \in \mathrm{tp}(c/Ma)$ and $\psi(z; b, m) \in \mathrm{tp}(c/Mb)$ which isolate these types. In particular, for some $k, k' < \omega$, we may choose χ and ψ so that for all a', there are at most k realizations of $\chi(z; a', m)$, and for all b', there are at most k' realizations of $\psi(z; b', m)$. Note that, since $\psi(z; b, m)$ isolates $\mathrm{tp}(c/Mb)$, if $b' \equiv_M b$, then none of the realizations of $\psi(z; b', m)$ are in M.

Let $\varphi(x, b, m)$ be the formula $\exists z\,(\chi(z; x, m) \wedge \psi(z; b, m))$. For any a' satisfying $\varphi(x, b, m)$, the witness to the existential quantifier also witnesses $a' \not\mathop{\smile\!\!\!\!\!|}\, [a]_M b$. This verifies strong finite character.

We use the same formula $\varphi(x, b, m)$ for witnessing. Let $(b_i)_{i<\omega}$ be a Morley sequence over M in a global M-finitely satisfiable type extending $\mathrm{tp}(b/M)$. If $\mathbb{M} \models \exists z\,(\psi(z; b_i, m) \wedge \psi(z; b_j, m))$ with $i < j$, then by finite satisfiability, there exists $m' \in M$ such that $\mathbb{M} \models \exists z\,(\psi(z; b_i, m) \wedge \psi(z; m', m))$. But every realization of $\psi(z; m', m)$ is algebraic over M and hence in M, while no realization of $\psi(z; b_i, m)$ is in M. It follows that the sets $\{\psi(\mathbb{M}; b_i, m) \mid i < \omega\}$ are pairwise disjoint, and thus the partial type $\{\varphi(x, b_i, m) \mid i < \omega\}$ is inconsistent, since for any a', the set $\chi(\mathbb{M}; a', m)$ intersects at most k of the sets $\psi(\mathbb{M}; b_i, m)$. $\square$

In Section 6.4, we will need strengthenings of the extension property and the independence theorem, which tell us that Kim-independence interacts with algebraic independence in a reasonable way. Along the way to proving the strengthening of the independence theorem, we will need a similar strengthening of the chain condition.

Definition 6.2.8. Using the same notation as in Theorem 6.2.4, we define the following properties of an abstract independence relation $\mathop{\smile\!\!\!\!\!|}$:

- *Algebraically reasonable extension*: if $a \mathop{\smile\!\!\!\!\!|}_M b$, then for any c, there exists a' such that $a' \equiv_{Mb} a$ and $a' \mathop{\smile\!\!\!\!\!|}_M bc$, and further $a' \mathop{\smile\!\!\!\!\!|}^a_{Mb} c$.
- *The algebraically reasonable chain condition*: if $a \mathop{\smile\!\!\!\!\!|}_M b$ and $I = (b_i)_{i<\omega}$ is a Morley sequence over M in a global M-invariant type extending $\mathrm{tp}(b/M)$, then there exists a' such that $a' \equiv_{Mb} a$, $a' \mathop{\smile\!\!\!\!\!|}_M I$, and I is Ma'-indiscernible, and further $b_i \mathop{\smile\!\!\!\!\!|}^a_{Ma'} b_j$ for all $i \neq j$.
- *The algebraically reasonable independence theorem*: if $a \mathop{\smile\!\!\!\!\!|}_M b$, $a' \mathop{\smile\!\!\!\!\!|}_M c$, $b \mathop{\smile\!\!\!\!\!|}_M c$, and $a \equiv_M a'$, then there exists a'' such that $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$, and $a'' \mathop{\smile\!\!\!\!\!|}_M bc$, and further $a'' \mathop{\smile\!\!\!\!\!|}^a_{Mb} c$, $a'' \mathop{\smile\!\!\!\!\!|}^a_{Mc} b$, and $b \mathop{\smile\!\!\!\!\!|}^a_{Ma''} c$.

Remark 6.2.9. If T is simple, then by Proposition III.8.4, $\mathop{\smile\!\!\!\!\!|}^K$ coincides with forking independence $\mathop{\smile\!\!\!\!\!|}^f$ over models, and the "and further" clauses of Definition 6.2.8 follow easily from the basic properties of forking independence.

Remark 6.2.10. In *any* theory, forking independence satisfies base monotonicity and strengthens algebraic independence. So for any set A and tuples a, b, and c, $a \mathop{\smile\!\!\!\!\!|}^f_A bc$ implies $a \mathop{\smile\!\!\!\!\!|}^f_{Ab} c$, which implies $a \mathop{\smile\!\!\!\!\!|}^a_{Ab} c$. Even in an NSOP$_1$ theory T, however, it is possible to have a model $M \models T$ and tuples a, b, and c, with $a \mathop{\smile\!\!\!\!\!|}^K_M bc$ and $a \not\mathop{\smile\!\!\!\!\!|}^a_{Mb} c$. See Example 6.3.15 below.

In the remainder of this section, we will show that in an NSOP$_1$ theory, Kim-independence satisfies the algebraically reasonable properties in Definition 6.2.8. The reader who is not interested in the technicalities of these proofs may skip directly to Section 6.3.

An improved independence theorem

We will first establish a slight improvement to the conclusion of the independence theorem, removing the apparent asymmetry between a, b, and c in the conclusion. As in Remark 6.2.9, this improved statement is easy in the context of a simple theory, where it follows from the basic properties of forking independence.

Definition 6.2.11. Suppose T is NSOP$_1$, $M \prec \mathbb{M}$, and $(a_i)_{i<\omega}$ is an M-indiscernible sequence.

1. Say $(a_i)_{i<\omega}$ is a *witness* for Kim-dividing over M if, whenever $\varphi(x; a_0)$ Kim-divides over M, $\{\varphi(x; a_i) \mid i < \omega\}$ is inconsistent.

2. Say $(a_i)_{i<\omega}$ is a *strong witness* to Kim-dividing over M if, for all $n < \omega$, the sequence $(a_{n \cdot i}, a_{n \cdot i+1}, \ldots, a_{n \cdot i+n-1})_{i<\omega}$ is a witness to Kim-dividing over M.

3. If I is any ordered index set, we say $(a_i)_{i \in I}$ is a *strong witness* to Kim-dividing over M if it has the same EM-type as a strong witness to Kim-dividing over M indexed by ω.

By Proposition III.7.9, in an NSOP$_1$ theory, the class of strong witnesses to Kim-dividing over M coincides with the *tree Morley sequences* over M. As we will not need the Morley tree machinery of Chapter 3, we will refer only to strong witnesses. The following facts are all contained in Chapter 3:

Fact 6.2.12. Suppose T is NSOP$_1$ and $M \models T$.

1. Suppose $(a_i, b_i)_{i \in I}$ is a strong witness to Kim-dividing over M and $J \subseteq I$ is an infinite subset. Then $(a_i)_{i \in J}$ and $(b_i)_{i \in J}$ are strong witnesses to Kim-dividing over M.

2. If $b \equiv_M b'$ and $b \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M b'$, then there is a strong witness to Kim-dividing over M, $(b_i)_{i \in \mathbb{Z}}$, with $b_0 = b$ and $b_1 = b'$.

3. $a \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M b$ if and only if there is an Ma-indiscernible sequence $(b_i)_{i<\omega}$ which is a strong witness to Kim-dividing over M with $b_0 = b$.

4. If $a \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M b$, and $I = (b_i)_{i \in \mathbb{Z}}$ is a strong witness to Kim-dividing over M with $b_0 = b$, then there exists $a' \equiv_{Mb} a$ such that I is Ma'-indiscernible and $a' \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M I$.

Theorem 6.2.13. *Suppose T is* NSOP$_1$, *$M \prec \mathbb{M}$, $a_0 \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M b$, $a_1 \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M c$, $b \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M c$ and $a_0 \equiv_M a_1$. Then there exists a with $a \equiv_{Mb} a_0$, $a \equiv_{Mc} a_1$, and $a \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M bc$, and further $b \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M ac$ and $c \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M ab$.*

Proof. Applying the independence theorem, we obtain a_2 with $a_2 \equiv_{Mb} a_0$, $a_2 \equiv_{Mc} a_1$, and $a_2 \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M bc$. Since $b \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M c$, by extension and an automorphism, there exists b' with $b' \equiv_{Mc} b$ such that $b \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M b'c$. By symmetry, $b'c \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M b$, and by extension and an automorphism again, there exists c' with $c' \equiv_{Mb} c$ such that $b'c \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^K_M bc'$. Altogether, $b'c \equiv_M bc \equiv_M bc'$, so by

Fact 6.2.12(2), there is a strong witness to Kim-dividing over M, $I = (b_i, c_i)_{i\in\mathbb{Z}}$, so that $(b_0, c_0) = (b, c')$ and $(b_1, c_1) = (b', c)$.

Choose a_3 such that $a_3bc' \equiv_M a_2bc$. Then $a_3 \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M bc'$, so by Fact 6.2.12(4), there exists $a \equiv_{Mbc'} a_3$ such that I is Ma-indiscernible and $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M I$. By monotonicity, $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M bc$. And we have $a \equiv_{Mb} a_3 \equiv_{Mb} a_2 \equiv_{Mb} a_0$, and by indiscernibility, $ac \equiv_M ac' \equiv_M a_3c' \equiv_M a_2c \equiv_M a_1c$, so $a \equiv_{Mc} a_1$.

By Fact 6.2.12(1), $(b_i)_{i\leq 0}$ is a strong witness to Kim-dividing over M with $b_0 = b$ which is Mac-indiscernible, so $b \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M ac$ by symmetry and Fact 6.2.12(3). Likewise, $(c_i)_{i\geq 1}$ is a strong witness to Kim-dividing over M with $c_1 = c$ which is Mab-indiscernible, so $c \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M ab$. This completes the proof. □

As an immediate corollary of the strengthened independence theorem, we get the following form of extension.

Corollary 6.2.14. *Suppose T is* NSOP$_1$ *and $M \models T$. If $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$ and $c \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M a$, then there is $c' \equiv_{Ma} c$ such that $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M bc'$ and $c' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M ab$.*

Proof. By extension, choose $c_* \equiv_M c$ with $c_* \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$. Then by Theorem 6.2.13, there exists c' such that $c' \equiv_{Ma} c$, $c' \equiv_{Mb} c_*$, $c' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M ab$, $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M bc'$, and $b \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M ac'$, which is more than we need. □

Kim-independence and algebraic independence

We are now ready to show that Kim-independence satisfies the algebraically reasonable conditions of Definition 6.2.8 in any NSOP$_1$ theory.

Theorem 6.2.15. *If T is* NSOP$_1$*, then $\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K$ satisfies algebraically reasonable extension.*

Proof. Suppose we have a model M and tuples a, b, c, with $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M b$. Let $\overline{b}$ be a tuple enumerating $\mathrm{acl}(Mb)$, and let $\overline{c} = (c_i)_{i\in I}$ be a tuple enumerating $\mathrm{acl}(Mbc) \setminus \mathrm{acl}(Mb)$. Then we also have $a \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M \overline{b}$ (Corollary III.5.17). And for any a', we have $a' \not\mathop{\smile\!\!\!\!\!\!\mid}\nolimits^a_{M\overline{b}} \overline{c}$ if and only if there is some index $i \in I$ such that $c_i \in \mathrm{acl}(Ma'b)$.

Let $\kappa = |\,\mathrm{acl}(Mab)|$. Since $\overline{b}$ enumerates an algebraically closed set, we can find pairwise disjoint tuples $(\overline{c}^\alpha)_{\alpha<\kappa^+}$ such that $\overline{c}^\alpha \equiv_{M\overline{b}} \overline{c}$ for all α. By extension, there exists $a'' \equiv_{M\overline{b}} a$ such that $a'' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M \overline{b}(\overline{c}^\alpha)_{\alpha<\kappa^+}$. In particular, for every α, $a'' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M \overline{b}\overline{c}^\alpha$. And since $|\,\mathrm{acl}(Ma''b)| = \kappa$, there is some α such that $\overline{c}^\alpha$ is disjoint from $\mathrm{acl}(Ma''b)$, so $a'' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^a_{M\overline{b}} \overline{c}^\alpha$.

Let σ be an automorphism moving $\overline{c}^\alpha$ to $\overline{c}$ and fixing $\overline{b}$, and let $a' = \sigma(a'')$. Then $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M bc$ and $a' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^a_{Mb} c$. □

Lemma 6.2.16. *Suppose T is* NSOP$_1$ *and $M \models T$. If $b \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M a$ and $c \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M a$, then there is $c' \equiv_{Ma} c$ such that $bc' \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^K_M a$ and $b \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^a_{Ma} c'$.*

Proof. We build a sequence $(b_i)_{i<\omega}$ by induction, such that for all $i < \omega$, the following conditions hold:

1. $b_i \equiv_{Ma} b$.
2. $b_{i+1} \mathop{\smile\!\!\!\!\!\mid}^K_M ab_{\leq i}$.
3. $b_{i+1} \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} b_{\leq i}$.

Set $b_0 = b$, and given $b_{\leq i}$, as $b \mathop{\smile\!\!\!\!\!\mid}^K_M a$, use algebraically reasonable extension to find $b_{i+1} \equiv_{Ma} b$ such that $b_{i+1} \mathop{\smile\!\!\!\!\!\mid}^K_M ab_{\leq i}$ and $b_{i+1} \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} b_{\leq i}$.

In particular, note that $b_i \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} b_j$ for all $i \neq j$.

Define a partial type $\Gamma(x; a, b)$ by

$$\Gamma(x; a, b) = \mathrm{tp}(c/Ma) \cup \{\neg\varphi(x, b; a) \mid \varphi(x, y; a) \text{ Kim-divides over } M\}.$$

Claim: For all $n < \omega$, $\bigcup_{i \leq n} \Gamma(x; a, b_i)$ is consistent.

Proof of claim: By induction on n, we will find $c_n \mathop{\smile\!\!\!\!\!\mid}^K_M ab_{\leq n}$ so that $c_n \models \bigcup_{i\leq n} \Gamma(x; a, b_i)$, i.e. $c_n \equiv_{Ma} c$ and $b_i c \mathop{\smile\!\!\!\!\!\mid}^K_M a$ for all $i \leq n$.

For $n = 0$, the existence of such a c_0 is given by Corollary 6.2.14. Suppose we have $c_n \mathop{\smile\!\!\!\!\!\mid}^K_M ab_{\leq n}$ realizing $\bigcup_{i\leq n} \Gamma(x; a, b_i)$. By extension, choose $c' \equiv_M c$ with $c' \mathop{\smile\!\!\!\!\!\mid}^K_M b_{n+1}$. As $b_{n+1} \mathop{\smile\!\!\!\!\!\mid}^K_M ab_{\leq n}$, we may apply the strengthened independence theorem (Theorem 6.2.13), to find $c_{n+1} \models \mathrm{tp}(c_n/Mab_{\leq n}) \cup \mathrm{tp}(c'/Mb_{n+1})$ with $c_{n+1} \mathop{\smile\!\!\!\!\!\mid}^K_M ab_{\leq n+1}$ and $b_{n+1}c_{n+1} \mathop{\smile\!\!\!\!\!\mid}^K_M ab_{\leq n}$. In particular, $b_{n+1}c_{n+1} \mathop{\smile\!\!\!\!\!\mid}^K_M a$, so $c_{n+1} \models \Gamma(x; a, b_{n+1})$. This gives $c_{n+1} \models \bigcup_{i\leq n+1} \Gamma(x; a, b_i)$. □

Let $\kappa = |\operatorname{acl}(Mac)|$, and let $(b'_\alpha)_{\alpha<\kappa^+}$ be an Ma-indiscernible sequence locally based on $(b_i)_{i<\omega}$. Then we have $b'_\alpha \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} b'_\beta$ for all $\alpha \neq \beta$, and $\bigcup_{\alpha<\kappa^+} \Gamma(x; a, b'_i)$ is consistent, by the claim. Let c_* realize this partial type, so $c_* \equiv_{Ma} c$ and $b'_\alpha c_* \mathop{\smile\!\!\!\!\!\mid}^K_M a$ for all α.

Since the sets $\operatorname{acl}(Mab'_\alpha)$ are pairwise disjoint over $\operatorname{acl}(Ma)$, and $|\operatorname{acl}(Mac_*)| = \kappa$, there is some $\alpha < \kappa^+$ such that $\operatorname{acl}(Mab'_\alpha)$ is disjoint from $\operatorname{acl}(Mac)$ over $\operatorname{acl}(Ma)$, so $b'_\alpha \mathop{\smile\!\!\!\!\!\mid}^a_M c$. Since $b_\alpha \equiv_{Ma} b$, we can find an automorphism σ fixing Ma and moving b'_α to b. Taking $c' = \sigma(c_*)$, we have $c' \equiv_{Ma} c_* \equiv_{Ma} c$ and $bc' \mathop{\smile\!\!\!\!\!\mid}^K_M a$, as desired. □

Corollary 6.2.17. *Suppose T is* NSOP$_1$ *and $M \models T$. If $b \mathop{\smile\!\!\!\!\!\mid}^K_M a$, then for any cardinal κ, there is an Ma-indiscernible sequence $I = (b_\alpha)_{\alpha<\kappa}$ with $b_0 = b$, such that $b_\alpha \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} b_\beta$ for all $\alpha \neq \beta$, and $I \mathop{\smile\!\!\!\!\!\mid}^K_M a$.*

Proof. We first build a sequence $(c_i)_{i<\omega}$ by induction, such that for all $i < \omega$, the following conditions hold:

1. $c_i \equiv_{Ma} b$.
2. $c_i \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} c_{<i}$.

3. $c_{\leq i} \mathop{\smile\!\!\!\!\!\mid}^K_M a$.

Set $c_0 = b$, and given $c_{\leq i}$, as $c_{\leq i} \mathop{\smile\!\!\!\!\!\mid}^K_M a$ and $b \mathop{\smile\!\!\!\!\!\mid}^K_M a$, we may apply Lemma 6.2.16 to find $c_{i+1} \equiv_{Ma} b$ such that $c_{i+1} \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} c_{\leq i}$ and $c_{\leq i+1} \mathop{\smile\!\!\!\!\!\mid}^K_M a$.

Now let $I = (b_\alpha)_{\alpha<\kappa}$ be an Ma-indiscernible sequence locally based on $(c_i)_{i<\omega}$. By condition (1), we may assume that $b_0 = b$. Condition (2) implies that $c_i \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} c_j$ for all $i \neq j$, so $b_\alpha \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} b_\beta$ for all $\alpha \neq \beta$, and $I \mathop{\smile\!\!\!\!\!\mid}^K_M a$ by condition (3) and the strong finite character of Kim-dividing. □

Theorem 6.2.18. *If T is* NSOP$_1$*, then $\mathop{\smile\!\!\!\!\!\mid}^K$ satisfies the algebraically reasonable chain condition.*

Proof. Suppose $a \mathop{\smile\!\!\!\!\!\mid}^K_M b$, and let $I = (b_i)_{i\in\omega}$ be a Morley sequence over M in a global M-invariant type $q \supseteq \mathrm{tp}(b/M)$.

Claim: For all n, there exists $(c_0, \dots, c_n) \models q^{\otimes(n+1)}|_M$ such that $c_i \equiv_{Ma} b$ for all $i \leq n$, $c_i \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} c_j$ for all $i \neq j$, and $(c_i)_{i\leq n} \mathop{\smile\!\!\!\!\!\mid}^K_M a$.

Proof of claim: By induction on n. When $n = 0$, taking $c_0 = b$ suffices. So suppose we are given the tuple $(c_0, \dots, c_n)$ by induction. Let $\kappa = |\operatorname{acl}(Mab)|$, and, applying Corollary 6.2.17, let $J = (d_{0,\alpha})_{\alpha<\kappa^+}$ be an Ma-indiscernible sequence with $d_{0,0} = b$, such that $d_{0,\alpha} \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} d_{0,\beta}$ for all $\alpha \neq \beta$, and $J \mathop{\smile\!\!\!\!\!\mid}^K_M a$.

Let $(d_1, \dots, d_{n+1})$ realize $q^{\otimes(n+1)}|_{MJ}$. Since $d_{0,\alpha} \models q|_M$ for all α, we have $(d_{0,\alpha}, d_1, \dots, d_{n+1}) \models q^{\otimes(n+2)}|_M$ for all α. Let $\sigma \in \mathrm{Aut}(\mathbb{M}/M)$ be such that $\sigma(c_i) = d_{i+1}$ for all $i \leq n$, and let $a' = \sigma(a)$. Now $a \equiv_M a'$, $a \mathop{\smile\!\!\!\!\!\mid}^K_M J$ (by choice of J), $a' \mathop{\smile\!\!\!\!\!\mid}^K_M (d_1, \dots, d_{n+1})$ (by invariance), and $(d_1, \dots, d_{n+1}) \mathop{\smile\!\!\!\!\!\mid}^K_M J$ (since $\mathrm{tp}(d_1, \dots, d_{n+1}/MJ)$ extends to a global M-invariant type).

Applying the independence theorem, we find a'' with $a'' \equiv_{MJ} a$, $a'' \equiv_{Md_1\dots d_{n+1}} a'$, and $a'' \mathop{\smile\!\!\!\!\!\mid}^K_M Jd_1 \dots d_{n+1}$. Then we still have $d_i \mathop{\smile\!\!\!\!\!\mid}^a_{Ma''} d_j$ for all $1 \leq i < j \leq n+1$, and since the sets $\operatorname{acl}(Ma''d_{0,\alpha})$ are pairwise disjoint over $\operatorname{acl}(Ma'')$, and $|\operatorname{acl}(Ma''d_i)| = \kappa$ for all $1 \leq i \leq n$, there is some $\alpha < \kappa^+$ such that $\operatorname{acl}(Ma''d_{0,\alpha})$ is disjoint from each of these n sets over $\operatorname{acl}(Ma'')$. So setting $d_0 = d_{0,\alpha}$, we have $d_i \mathop{\smile\!\!\!\!\!\mid}^a_{Ma''} d_j$ for all $i \neq j$.

It remains to move a'' back to a by an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}/M)$. The tuple $\sigma(d_0, \dots, d_{n+1})$ has the desired properties. □

By compactness, we can find $I' = (c_i)_{i<\omega} \models q^{\otimes\omega}$, i.e. a q-Morley sequence over M, such that $c_i \equiv_{Ma} b$ for all $i < \omega$, $c_i \mathop{\smile\!\!\!\!\!\mid}^a_{Ma} c_j$ for all $i \neq j$, and $I' \mathop{\smile\!\!\!\!\!\mid}^K_M a$. In fact, we can assume that I' is Ma-indiscernible, by replacing it with an Ma-indiscernible sequence locally based on it. As I' and I are both q-Morley sequences over M, we can move I' to I by an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}/M)$, and take $a' = \sigma(a)$. □

Lemma 6.2.19. *Suppose $(a_i)_{i<\omega}$ is a Morley sequence for a global M-invariant type, which is moreover Mb-indiscernible. If $b \mathop{\smile\!\!\!\!\!\mid}^a_{M(a_i)_{i<\omega}} b'$, then $b \mathop{\smile\!\!\!\!\!\mid}^a_{Ma_0} b'$.*

Proof. Suppose there is some element $c \in \operatorname{acl}(Ma_0b) \cap \operatorname{acl}(Ma_0b')$. We would like to show that $c \in \operatorname{acl}(Ma_0)$. What we have is that $c \in \operatorname{acl}(M(a_i)_{i<\omega})$, and in particular $c \in \operatorname{acl}(Ma_0b) \cap$

$\operatorname{acl}(Ma_0a_{i_1}\dots a_{i_n})$ for some $0 < i_1 < \dots < i_n$. Now $(a_i)_{i\geq 1}$ is indiscernible over Ma_0b, hence indiscernible over $\operatorname{acl}(Ma_0b)$, which contains c. So c is also in both $\operatorname{acl}(Ma_0a_1\dots a_n)$ and $\operatorname{acl}(Ma_0a_{n+1}\dots a_{2n})$.

But $\operatorname{tp}(a_{n+1},\dots,a_{2n}/Ma_0a_1\dots a_n)$ extends to a global Ma_0-invariant type, so we must have $a_{n+1}\dots a_{2n} \mathop{\smile\!\!\!\!\!\!|}^a_{Ma_0} a_1\dots a_n$. Hence $c\in\operatorname{acl}(Ma_0)$. □

Lemma 6.2.20. *Given $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$ and global M-invariant types $p(x)$ and $q(y)$ extending $\operatorname{tp}(a/M)$ and $\operatorname{tp}(b/M)$ respectively, there exist mutually indiscernible Morley sequences $(a_i)_{i<\omega}$ and $(b_i)_{i<\omega}$ in p and q, with $a_0 = a$ and $b_0 = b$, such that $(a_i)_{i<\omega} \mathop{\smile\!\!\!\!\!\!|}^K_M (b_i)_{i<\omega}$, and $a_i \mathop{\smile\!\!\!\!\!\!|}^a_{Mb} a_j$ and $b_i \mathop{\smile\!\!\!\!\!\!|}^a_{Ma} b_j$ for all $i\neq j$.*

Proof. Let $(a_i)_{i<\omega}$ be a Morley sequence in p, with $a_0 = a$. By the algebraically reasonable chain condition, there is some $b' \equiv_{Ma} b$, such that $(a_i)_{i<\omega}$ is Mb'-indiscernible, $b' \mathop{\smile\!\!\!\!\!\!|}^K_M (a_i)_{i<\omega}$, and $a_i \mathop{\smile\!\!\!\!\!\!|}^a_{Mb'} a_j$ for all $i\neq j$. At the expense of moving $(a_i)_{i<\omega}$ by an automorphism fixing Ma, we may assume that $b' = b$.

Now let $(b_i)_{i<\omega}$ be a Morley sequence in q, with $b_0 = b$. Since $(a_i)_{i<\omega} \mathop{\smile\!\!\!\!\!\!|}^K_M b$, we can find some $(a'_i)_{i<\omega} \equiv_{Mb} (a_i)_{i<\omega}$, such that $(b_i)_{i<\omega}$ is $M(a'_i)_{i<\omega}$-indiscernible, $(a'_i)_{i<\omega} \mathop{\smile\!\!\!\!\!\!|}^K_M (b_i)_{i<\omega}$, and $b_i \mathop{\smile\!\!\!\!\!\!|}^a_{M(a'_i)_{i<\omega}} b_j$ for all $i\neq j$. Further, we may replace $(a'_i)_{i<\omega}$ with an $M(b_i)_{i<\omega}$-indiscernible sequence $(a''_i)_{i<\omega}$ locally based on it, and at the expense of moving $(b_i)_{i<\omega}$ by an automorphism fixing Mb, we may assume that $(a''_i)_{i<\omega} = (a_i)_{i<\omega}$.

The result is that $(a_i)_{i<\omega}$ and $(b_i)_{i<\omega}$ are mutually indiscernible Morley sequences in p and q with $a_0 = a$ and $b_0 = b$ and $(a_i)_{i<\omega} \mathop{\smile\!\!\!\!\!\!|}^K_M (b_i)_{i<\omega}$. We have also ensured that $a_i \mathop{\smile\!\!\!\!\!\!|}^a_{Mb} a_j$ for all $i\neq j$ and $b_i \mathop{\smile\!\!\!\!\!\!|}^a_{M(a_i)_{i<\omega}} b_j$ for all $i\neq j$. By Lemma 6.2.19, also $b_i \mathop{\smile\!\!\!\!\!\!|}^a_{Ma} b_j$ for all $i\neq j$. □

Theorem 6.2.21. *If T is NSOP_1, then $\mathop{\smile\!\!\!\!\!\!|}^K$ satisfies the algebraically reasonable independence theorem.*

Proof. We have a model M and tuples a, a', b, c, with $a \mathop{\smile\!\!\!\!\!\!|}^K_M b$, $a' \mathop{\smile\!\!\!\!\!\!|}^K_M c$, $b \mathop{\smile\!\!\!\!\!\!|}^K_M c$, and $a \equiv_M a'$. Let $p(x)$, $q(y)$, and $r(z)$ be global M-invariant types extending $\operatorname{tp}(a/M) = \operatorname{tp}(a'/M)$, $\operatorname{tp}(b/M)$, and $\operatorname{tp}(c/M)$, respectively.

Apply Lemma 6.2.20 to $q(y)$ and $r(z)$, obtaining Morley sequences $(b_i)_{i<\omega}$ and $(c_i)_{i<\omega}$. Then apply it two more times, to $p(x)$ and $q(y)$, obtaining Morley sequences $(a_i)_{i<\omega}$ and $(\widehat{b}_i)_{i<\omega}$, and then to $p(x)$ and $r(z)$, obtaining Morley sequences $(a'_i)_{i<\omega}$ and $(\widehat{c}_i)_{i<\omega}$. At the expense of moving $(a_i)_{i<\omega}$ and $(a'_i)_{i<\omega}$ by automorphisms over M, we may assume that $(\widehat{b}_i)_{i<\omega} = (b_i)_{i<\omega}$ and $(\widehat{c}_i)_{i<\omega} = (c_i)_{i<\omega}$. Note that $(a_i)_{i<\omega}$ and $(a'_i)_{i<\omega}$ are both p-Morley sequences over M, so $(a_i)_{i<\omega} \equiv_M (a'_i)_{i<\omega}$.

We now apply the independence theorem to the sequences $(a_i)_{i<\omega}$, $(a'_i)_{i<\omega}$, $(b_i)_{i<\omega}$, and $(c_i)_{i<\omega}$, obtaining a sequence $(a''_i)_{i<\omega}$ such that $(a''_i)_{i<\omega} \equiv_{M(b_i)_{i<\omega}} (a_i)_{i<\omega}$, $(a''_i)_{i<\omega} \equiv_{M(c_i)_{i<\omega}} (a'_i)_{i<\omega}$, and $(a''_i)_{i<\omega} \mathop{\smile\!\!\!\!\!\!|}^K_M (b_i)_{i<\omega}(c_i)_{i<\omega}$. The sequences $(a''_i)_{i<\omega}$, $(b_i)_{i<\omega}$, and $(c_i)_{i<\omega}$ are pairwise mutually indiscernible over M and have the property that any pair from one sequence is algebraically independent over any element of another sequence.

Let κ be a cardinal larger than the sizes of M, the language, and the tuples a, b, and c. We can stretch the (b_i) sequence to have length κ^+ and stretch the (c_i) sequence to have length κ^{++}, while maintaining their mutual indiscernibility and algebraic independence properties.

Fix $a_* = a_0''$. For any $i < \kappa^+$, $|\operatorname{acl}(Ma_*b_i)| \leq \kappa$. Since the sets $\{\operatorname{acl}(Ma_*c_j) \mid j < \kappa^{++}\}$ are pairwise disjoint over $\operatorname{acl}(Ma_*)$, we can remove any c_j such that there exists an $i < \kappa^+$ such that $b_i \not\mathop{\smile}\limits^{\mid}{}^{[a]}_{Ma_*} c_j$, and we are still left with a sequence of length κ^{++}.

Similarly, since for all $i < \kappa^+$, the sets $\{\operatorname{acl}(Mb_ic_j) \mid j < \kappa^{++}\}$ are pairwise disjoint over $\operatorname{acl}(Mb_i)$, we can further remove any c_j such that there exists an i such that $a \not\mathop{\smile}\limits^{\mid}{}^{[a]}_{Mb_i} c_j$, and we are still left with a sequence of length κ^{++}. Fix a c_* from this sequence.

Finally, since the sets $\{\operatorname{acl}(Mb_ic_*) \mid i < \kappa^+\}$ are pairwise disjoint over $\operatorname{acl}(Mc_*)$, we can remove any b_i such that $a_* \not\mathop{\smile}\limits^{\mid}{}^{[a]}_{Mc_*} b_i$, and we are still left with a sequence of length κ^+. Fix a b_* from this sequence.

It remains to move b_*c_* back to bc by an automorphism σ fixing M, and set $a'' = \sigma(a_*)$. □

6.3 The model companion of the empty theory

The theory $T^{\emptyset}_L$

Let L be any language. Then the empty L-theory has a model completion, which we call $T^{\emptyset}_L$. As the theory $T^{\emptyset}_L$ may be regarded as the generic expansion of the theory of an infinite set in the empty language, this fact is a special case of Theorem 5 in [Win75] (see Fact 6.4.2 below), and it was reproven by Jeřábek in [Jeř17]. We include a proof, following the idea of [Win75], for completeness and to fix notation.

Definition 6.3.1. A *partial diagram* Δ is a set of atomic and negated atomic formulas. Δ is *flat* if each formula in Δ has the form $R(\overline{z})$, $\neg R(\overline{z})$, or $f(\overline{z}) = z'$, where $\overline{z}$ is a tuple of variables and z' is a single variable. We view constant symbols as 0-ary function symbols, so this includes formulas of the form $c = z'$.

In a flat diagram, we always intend distinct variables to refer to distinct elements.

Definition 6.3.2. A flat diagram Δ is *consistent* if, for each tuple of variables $\overline{z}$,

1. At most one of $R(\overline{z})$ and $\neg R(\overline{z})$ is in Δ, where R is a relation symbol.
2. There is at most one variable z' such that $f(\overline{z}) = z'$ is in Δ, where f is a function symbol.

Definition 6.3.3. A consistent flat diagram Δ in the variables $\overline{w}$ is *complete* if, for each tuple of variables $\overline{z}$ from $\overline{w}$,

1. Either $R(\overline{z})$ or $\neg R(\overline{z})$ is in Δ.

2. There is some variable z' in $\overline{w}$ such that $f(\overline{z}) = z'$ is in Δ.

Let A be any L-structure. Then there is a complete flat diagram $\mathrm{diag}_f(A)$ in the variables $(w_a)_{a\in A}$, which contains a formula $\psi(w_{a_1}, \ldots, w_{a_n})$ of one of the allowed forms if and only if $A \models \psi(a_1, \ldots, a_n)$. The following easy lemma establishes the converse.

Lemma 6.3.4. *Suppose Δ is a consistent flat diagram in the nonempty set of variables $(w_a)_{a\in A}$. Then there is an L-structure with domain A such that for all $\psi(w_{a_1}, \ldots, w_{a_n}) \in \Delta$, $A \models \psi(a_1 \ldots, a_n)$.*

Proof. First, we extend Δ to a complete flat diagram Δ' as follows: For each n-ary relation symbol R, and for each n-tuple $\overline{z}$ such that neither $R(\overline{z})$ nor $\neg R(\overline{z})$ is in Δ, add $\neg R(\overline{z})$ to Δ'. Now fix an arbitrary variable w_a. For each n-ary function symbol f, and for each n-tuple $\overline{z}$ such that no formula of the form $f(\overline{z}) = z'$ is in Δ, add $f(\overline{z}) = w_a$ to Δ'.

We define an L-structure with domain A, according to Δ'. If R is an n-ary relation symbol, we set $R^A = \{(a_1, \ldots, a_n) \in A^n \mid R(w_{a_1}, \ldots, w_{a_n}) \in \Delta'\}$. If f is an n-ary function symbol and $(a_1, \ldots, a_n) \in A^n$, we set $f^A(a_1, \ldots, a_n) = a'$, where a' is the unique element of A such that $f(w_{a_1}, \ldots, w_{a_n}) = w_{a'} \in \Delta'$. Consistency ensures that this L-structure is well-defined and satisfies all the formulas in Δ' (and hence in Δ). □

For the purposes of axiomatizing the existentially closed L-structures, we will be interested in a class of finite partial diagrams, which we call extension diagrams.

Definition 6.3.5. Let $\overline{w}$ be a finite tuple of variables, partitioned into two subtuples $\overline{x}$ and $\overline{y}$. An *extension diagram* in $(\overline{x}, \overline{y})$ is a consistent flat diagram Δ in the variables $\overline{w}$, such that for each formula $R(\overline{z})$, $\neg R(\overline{z})$, or $f(\overline{z}) = z'$ in Δ, some variable in $\overline{z}$ is in $\overline{y}$. In particular, no constant symbols appear in extension diagrams.

A tuple $\overline{a} = (a_i)_{i\in I}$ is *non-redundant* if $a_i \neq a_j$ for all $i \neq j$. Given a finite tuple of variables $\overline{z} = (z_1, \ldots, z_n)$, let $\delta(\overline{z})$ be the formula which says that $\overline{z}$ is non-redundant:

$$\bigwedge_{1\le i<j\le n} z_i \neq z_j.$$

Given a finite partial diagram Δ in the finite tuple of variables $\overline{w}$, let $\varphi_\Delta(\overline{w})$ be the conjunction of all the formulas in Δ, together with $\delta(\overline{w})$:

$$\left(\bigwedge_{\psi(\overline{z})\in\Delta} \psi(\overline{z})\right) \wedge \delta(\overline{w}).$$

Lemma 6.3.6. *Let Δ be an extension diagram in $(\overline{x}, \overline{y})$, and let A be an L-structure. If $\overline{a}$ is a non-redundant tuple from A of the same length as $\overline{x}$, then there is an L-structure B containing A and a tuple $\overline{b}$ from B of the same length as $\overline{y}$ such that $B \models \varphi_\Delta(\overline{a}, \overline{b})$.*

Proof. Consider the flat diagram $\text{diag}_f(A) \cup \Delta(x_{a_1}, \ldots, x_{a_n}, y_{b_1}, \ldots, y_{b_n})$, where we identify the variables $\overline{x}$ in Δ with the variables in $\text{diag}_f(A)$ enumerating $\overline{a}$, and we index the variables $\overline{y}$ in Δ by a new tuple $\overline{b}$. This diagram is consistent, since $\text{diag}_f(A)$ and Δ are individually consistent, and for every formula $R(\overline{z})$, $\neg R(\overline{z})$, or $f(\overline{z}) = z'$ in Δ, some element of the tuple $\overline{z}$ is in $\overline{y}$, while $\text{diag}_f(A)$ does not mention the variables in $\overline{y}$. Hence, by Lemma 6.3.4, there is a structure B with domain $A \cup \{b_1, \ldots, b_n\}$, such that B satisfies $\text{diag}_f(A)$ (so A is a substructure of B), and $B \models \varphi_\Delta(\overline{a}, \overline{b})$. □

Lemma 6.3.7. *If an L-structure A is not existentially closed, then there is a non-redundant tuple $\overline{a}$ from A and an extension diagram Δ in $(\overline{x}, \overline{y})$, such that $A \models \neg\exists\overline{y}\, \varphi_\Delta(\overline{a}, \overline{y})$.*

Proof. Since A is not existentially closed, there is a quantifier-free L-formula $\varphi(\overline{x}, \overline{y})$, an L-structure B containing A, and tuples $\overline{a} \in A$ and $\overline{b} \in B$, such that $B \models \varphi(\overline{a}, \overline{b})$, but $A \models \neg\exists\overline{y}\, \varphi(\overline{a}, \overline{y})$. We may assume the the tuples $\overline{a}$ and $\overline{b}$ are non-redundant and that $b_i \in B \setminus A$ for all i. Writing φ in disjunctive normal form, one of the disjuncts is satisfied by $(\overline{a}, \overline{b})$ in B, so we may assume that φ is a conjunction of atomic and negated atomic formulas. Let Δ be the finite partial diagram containing these formulas. Then $\varphi_\Delta(\overline{x}, \overline{y})$ is equivalent to $\varphi(\overline{x}, \overline{y}) \wedge \delta(\overline{x}, \overline{y})$, and we have $B \models \varphi_\Delta(\overline{a}, \overline{b})$, but $A \models \neg\exists\overline{y}\, \varphi_\Delta(\overline{a}, \overline{y})$.

We will transform Δ into an extension diagram. This process will involve adding and deleting variables and making corresponding changes to the tuples $\overline{a}$ and $\overline{b}$, but we will maintain the invariants that Δ is finite, $A \models \neg\exists\overline{y}\, \varphi_\Delta(\overline{a}, \overline{y})$, $B \models \varphi_\Delta(\overline{a}, \overline{b})$, and $b_i \in B \setminus A$ for all i. We write $\overline{w}$ for the tuple $(\overline{x}, \overline{y})$ and $\overline{c}$ for $(\overline{a}, \overline{b})$.

First, we flatten Δ. Suppose that there is an n-ary function symbol f such that the term $f(w_{i_1}, \ldots, w_{i_k})$ (where the w_{i_j} are variables) appears in a formula in Δ which is not of the form $f(w_{i_1}, \ldots, w_{i_k}) = w'$ for some variable w'. Let $d = f^B(c_{i_1}, \ldots, c_{i_k})$. If $d = c_{i_{k+1}}$ for some i_{k+1}, then we simply replace this instance of $f(w_{i_1}, \ldots, w_{i_k})$ with $w_{i_{k+1}}$ and add the formula $f(w_{i_1}, \ldots, w_{i_k}) = w_{i_{k+1}}$ to Δ if it is not already there. If d is not in the tuple $\overline{c}$, we introduce a new variable w' (a new x if $d \in A$ and a new y otherwise), add d to $\overline{c}$ (as a new a if $d \in A$ and a new b otherwise), replace this instance of $f(w_{i_1}, \ldots, w_{i_k})$ with w', and add the formula $f(w_{i_1}, \ldots, w_{i_k}) = w'$ to Δ.

Repeating this procedure, we eventually ensure that every formula in Δ has the form $w = w'$, $w \neq w'$, $R(w_{i_1}, \ldots, w_{i_n})$, $\neg R(w_{i_1}, \ldots, w_{i_n})$, or $f(w_{i_1}, \ldots, w_{i_n}) = w'$.

Next we remove the equations and inequations between variables. Since the tuples $\overline{a}$ and $\overline{b}$ are non-redundant, Δ does not contain any equalities between distinct variables, and the equalities of the form $w = w$ can of course be removed. Further, we may assume that Δ does not contain any inequalities $w_i \neq w_j$ between variables either, since these inequalities are all implied by $\delta(\overline{x}, \overline{y})$ and hence by φ_Δ. The set of formulas Δ is now a flat diagram. It is consistent, since it is satisfied by the non-redundant tuple $\overline{c}$.

Finally, let Δ' be the extension diagram obtained by removing from Δ any formula $R(\overline{z})$, $\neg R(\overline{z})$, or $f(\overline{z}) = z'$ in which none of the variables in $\overline{z}$ are in $\overline{y}$. Note that in the case of $f(\overline{z}) = z'$, if all of the $\overline{z}$ are in $\overline{x}$, then their interpretations come from A, and since A is closed under the function symbols, z' is in $\overline{x}$ as well.

So $\varphi_\Delta(\overline{x},\overline{y})$ is equivalent to $\varphi_{\Delta'}(\overline{x},\overline{y}) \wedge \bigwedge_{j=1}^N \psi_j(\overline{x})$, where each ψ_j is atomic or negated atomic. But since

$$A \models \neg\exists\overline{y} \left(\varphi_{\Delta'}(\overline{a},\overline{y}) \wedge \bigwedge_{j=1}^N \psi_j(\overline{a})\right),$$

also $A \models \neg\exists\overline{y}\,\varphi_{\Delta'}(\overline{a},\overline{y})$, as was to be shown. □

Given an extension diagram Δ in $(\overline{x},\overline{y})$, let ψ_Δ be the sentence

$$\forall\overline{x}\,(\delta(\overline{x}) \to \exists\overline{y}\,\varphi_\Delta(\overline{x},\overline{y})),$$

and let $T_L^\emptyset = \{\psi_\Delta \mid \Delta \text{ is an extension diagram in } (\overline{x},\overline{y})\}$.

Theorem 6.3.8. *$T_L^\emptyset$ is the model companion of the empty L-theory.*

Proof. It suffices to show that the class of existentially closed L-structures is axiomatized by $T_L^\emptyset$ [CK90, Proposition 3.5.15].

Suppose A is an existentially closed L-structure, and let Δ be an extension diagram in $(\overline{x},\overline{y})$. Let $\overline{a}$ be any non-redundant tuple from A of the same length as $\overline{x}$. By Lemma 6.3.6, there is an L-structure B containing A and a tuple $\overline{b}$ from B such that $B \models \varphi_\Delta(\overline{a},\overline{b})$. So $B \models \exists\overline{y}\,\varphi_\Delta(\overline{a},\overline{y})$. But since A is existentially closed, also $A \models \exists\overline{y}\,\varphi_\Delta(\overline{a},\overline{y})$. Hence $A \models \psi_\Delta$, and since Δ was arbitrary, $A \models T_L^\emptyset$.

Conversely, suppose the L-structure A is not existentially closed. By Lemma 6.3.7, there is a non-redundant tuple $\overline{a}$ from A and an extension diagram Δ in $(\overline{x},\overline{y})$ such that $A \models \neg\exists\overline{y}\,\varphi_\Delta(\overline{a},\overline{y})$. Hence $A \not\models \psi_\Delta$, and $A \not\models T_L^\emptyset$. □

Lemma 6.3.9. *The class of L-structures satisfies the disjoint amalgamation property.*

Proof. Let $f_1\colon A \to B$ and $f_2\colon A \to C$ be embeddings of L-structures. Let $B' = B \setminus f_1(A)$ and $C' = C \setminus f_2(A)$, and consider the diagrams $\mathrm{diag}_f(B) = \Delta_B((x_a)_{a\in A},(x_b)_{b\in B'})$ and $\mathrm{diag}_f(C) = \Delta_C((x_a)_{a\in A},(x_c)_{c\in C'})$, where we use the same variables $(x_a)_{a\in A}$ to enumerate $f_1(A)$ and $f_2(A)$.

Then $\Delta_B \cup \Delta_C$ is consistent, since the two diagrams agree on $\mathrm{diag}_f(A)$. By Lemma 6.3.4, we get an L-structure D with domain $A \cup B' \cup C'$, and the obvious maps $g_1\colon B \to D$ and $g_2\colon C \to D$ satisfy $g_1 \circ f_1 = g_2 \circ f_2$. These maps are embeddings, since D satisfies Δ_B and Δ_C, and the images of B and C are disjoint over the image of A. □

The following corollary now follows from standard facts about model completions (see [CK90, Proposition 3.5.18]).

Corollary 6.3.10. *$T_L^\emptyset$ is a model completion of the empty L-theory, and it has quantifier elimination. The completions of $T_L^\emptyset$ are obtained by specifying (by quantifier-free sentences) the isomorphism type of the structure $\langle\emptyset\rangle$ generated by the constants. Such a completion $\widetilde{T}$ is the model completion of the theory of L-structures containing a substructure isomorphic to $\langle\emptyset\rangle$. If there are no constant symbols in L, then $T_L^\emptyset$ is complete.*

Corollary 6.3.11. *Let $\mathbb{M}$ be a monster model for some completion of $T_L^\emptyset$. For any set $A \subseteq \mathbb{M}$, $\operatorname{acl}(A) = \operatorname{dcl}(A) = \langle A \rangle$.*

Proof. Since $\langle A \rangle \subseteq \operatorname{dcl}(A) \subseteq \operatorname{acl}(A)$, it suffices to show that $\operatorname{acl}(A) \subseteq \langle A \rangle$. Suppose $\varphi(\overline{a}, y)$ is an algebraic formula with parameters $\overline{a}$ from A, which is satisfied by exactly k elements of $\mathbb{M}$, including b. By Corollary 6.3.10, we may assume that φ is quantifier-free. Suppose for contradiction that $b \notin \langle A \rangle$, so that $\langle A \rangle$ is a proper substructure of $B = \langle Ab \rangle$. Let $C_0 = B$ and, by induction, apply Lemma 6.3.9 to obtain a disjoint amalgam C_{i+1} of C_i and B over $\langle A \rangle$. Let B_i denote the image of B in C_i. Then C_{k+1} contains $\langle A \rangle$, together with substructures $B_1, \dots, B_{k+1}$, pairwise disjoint over $\langle A \rangle$ and each isomorphic to B over $\langle A \rangle$. Then, by quantifier elimination and saturation, C_{k+1} embeds in $\mathbb{M}$ over $\langle A \rangle$, and we may identify the B_i with their images in $\mathbb{M}$. Each B_i contains an element b_i such that $\operatorname{qftp}(b_i/A) = \operatorname{qftp}(b/A)$, so $\mathbb{M} \models \varphi(\overline{a}, b_i)$ for all i, which is a contradiction. □

Independence and NSOP_1

For the remainder of this section, we fix a monster model $\mathbb{M} \models T_L^\emptyset$. As there is a monster model for every choice of completion of $T_L^\emptyset$ and $\mathbb{M}$ is arbitrary, to show that $T_L^\emptyset$ is NSOP_1, it suffices to establish this for $\operatorname{Th}(\mathbb{M})$.

Theorem 6.3.12. *$\mathop{\smile\!\!\!\!\!\!|}^a$ satisfies the independence theorem over arbitrary sets.*

Proof. Suppose we are given $C \subseteq \mathbb{M}$ and tuples a, a', b, c, with $a \mathop{\smile\!\!\!\!\!\!|}^a_C b$, $a' \mathop{\smile\!\!\!\!\!\!|}^a_C c$, $b \mathop{\smile\!\!\!\!\!\!|}^a_C c$, and $a \equiv_C a'$. Let x_C be a tuple enumerating $\langle C \rangle$, let x_a, x_b and x_c be tuples enumerating $\langle Ca \rangle \setminus \langle C \rangle$, $\langle Cb \rangle \setminus \langle C \rangle$, and $\langle Cc \rangle \setminus \langle C \rangle$, respectively, and let x_{ab}, x_{ac}, and x_{bc} be tuples enumerating $\langle Cab \rangle \setminus (\langle Ca \rangle \cup \langle Cb \rangle)$, $\langle Ca'c \rangle \setminus (\langle Ca' \rangle \cup \langle Cc \rangle)$, and $\langle Cbc \rangle \setminus (\langle Cb \rangle \cup \langle Mc \rangle)$, respectively.

Observe that (x_C, x_a, x_b, x_{ab}) enumerates $\langle Cab \rangle$ without repetitions. The only thing to check is that no elements of x_a and x_b name the same element of $\langle Cab \rangle$, and this is exactly the condition that $a \mathop{\smile\!\!\!\!\!\!|}^a_C b$. Similarly, (x_C, x_a, x_c, x_{ac}) enumerates $\langle Ca'c \rangle$ (where we view x_a as enumerating $\langle Ca' \rangle \setminus \langle C \rangle$ via the isomorphism $\langle Ca \rangle \to \langle Ca' \rangle$ induced by $a \mapsto a'$), and (x_C, x_b, x_c, x_{bc}) enumerates $\langle Cbc \rangle$.

Let $p_{ab} = \operatorname{diag}_f(\langle Cab \rangle)$, $p_{ac} = \operatorname{diag}_f(\langle Ca'c \rangle)$, and $p_{bc} = \operatorname{diag}_f(\langle Cbc \rangle)$. The flat diagram $p_{ab}(x_C, x_a, x_b, x_{ab}) \cup p_{ac}(x_C, x_a, x_c, x_{ac}) \cup p_{bc}(x_C, x_b, x_c, x_{bc})$ is consistent, since p_{ab}, p_{ac}, and p_{bc} agree on $\operatorname{diag}_f(\langle Ca \rangle) = \operatorname{diag}_f(\langle Ca' \rangle)$ (again, allowing x_a to enumerate $\langle Ca' \rangle \setminus \langle C \rangle$), $\operatorname{diag}_f(\langle Cb \rangle)$, and $\operatorname{diag}_f(\langle Cc \rangle)$. So by Lemma 6.3.4, it extends to the flat diagram of an L-structure X with domain $x_C \cup x_a \cup x_b \cup x_c \cup x_{ab} \cup x_{ac} \cup x_{bc}$.

Having constructed X, which agrees with $\mathbb{M}$ on the substructure generated by the empty set, we can embed it in $\mathbb{M}$ by $i \colon X \to \mathbb{M}$. Further,

$$\operatorname{qftp}(i(x_C), i(x_b), i(x_c), i(x_{bc})) = \operatorname{qftp}(\langle Cbc \rangle),$$

so by quantifier elimination

$$(i(x_C), i(x_b), i(x_c), i(x_{bc})) \equiv \langle Cbc \rangle,$$

and, by an automorphism of $\mathbb{M}$, we may assume that $i(x_C, x_b, x_c, x_{bc}) = \langle Cbc \rangle$. Let a'' the subtuple of $i(x_C, x_a)$ corresponding to the subtuple of (x_C, x_a) enumerating a.

Now $\mathrm{qftp}(i(x_C), i(x_a), i(x_b), i(x_{ab})) = \mathrm{qftp}(\langle Cab \rangle)$, so $a'' \equiv_{Cb} a$, and similarly

$$\mathrm{qftp}(i(x_C), i(x_a), i(x_c), i(x_{ac})) = \mathrm{qftp}(\langle Cac \rangle),$$

so $a'' \equiv_{Cc} a'$. Finally, $a'' \mathop{\smile\!\!\!\!\!\mid}\,^a_C bc$, since $i(x_a)$ enumerates $\langle Ca'' \rangle \setminus C$ and is disjoint from $(i(x_b), i(x_c), i(x_{bc}))$, which enumerates $\langle Cbc \rangle \setminus C$. □

Corollary 6.3.13. *$T^{\emptyset}_L$ is* NSOP$_1$ *and $\mathop{\smile\!\!\!\!\!\mid}\,^K = \mathop{\smile\!\!\!\!\!\mid}\,^a$ over models.*

Proof. By Theorem 6.3.12, Lemma 6.2.7, and Theorem 6.2.4. □

On the other hand, except in trivial cases, $T^{\emptyset}_L$ has TP$_2$ and therefore is not simple. For definitions of simple and TP$_2$ see, e.g., [Che14].

Proposition 6.3.14. *If L contains at least one n-ary function symbol with $n \geq 2$, then $T^{\emptyset}_L$ has TP$_2$, and is therefore not simple.*

Proof. In $\mathbb{M} \models T^{\emptyset}_L$, choose a set of pairwise distinct $(n-1)$-tuples $B = \{b_i : i < \omega\}$ and a set of pairwise distinct elements $C = \{c_{i,j} : i, j < \omega\}$ so that B and C are disjoint. Note that:

- For all $i < \omega$, $\{f(x, b_i) = c_{i,j} \mid j < \omega\}$ is 2-inconsistent.
- For all $g : \omega \to \omega$, $\{f(x, b_i) = c_{i,g(i)} \mid i < \omega\}$ is consistent.

Hence the formula $\varphi(x; y, z)$ given by $f(x, y) = z$ has TP$_2$, witnessed by the array $(b_i, c_{i,j})_{i<\omega, j<\omega}$. □

Forking and dividing

Next, we analyze forking and dividing in $T^{\emptyset}_L$. See [Adl05, Section 5] for the definitions of forking and dividing. We begin with an example of the distinction between forking independence and Kim-independence.

Example 6.3.15. Suppose L contains an n-ary function symbol f with $n \geq 2$. Let $M \models T^{\emptyset}_L$, let b be any $(n-1)$-tuple not in M, let c be any element not in $\langle Mb \rangle$, and let a be an element satisfying $a \mathop{\smile\!\!\!\!\!\mid}\,^a_M bc$, but $f(a, b) = c$. By Corollary 6.3.13, $a \mathop{\smile\!\!\!\!\!\mid}\,^K_M bc$. But $c \in (\langle Mab \rangle \cap \langle Mbc \rangle) \setminus \langle Mb \rangle$, so $a \not\mathop{\smile\!\!\!\!\!\mid}\,[a]_{Mb} c$. Then also $a \not\mathop{\smile\!\!\!\!\!\mid}\,[f]_M bc$, since otherwise we would have $a \mathop{\smile\!\!\!\!\!\mid}\,^f_{Mb} c$ and hence $a \mathop{\smile\!\!\!\!\!\mid}\,^a_{Mb} c$, by base monotonicity.

This example is closely related to the TP$_2$ array in Proposition 6.3.14. The formula $f(x, b) = c$ divides along any sequence $(bc_i)_{i<\omega}$ where b is constant but the c_i are distinct (like the rows of the TP$_2$ array). But this formula does not Kim-divide, since in any Morley sequence $(b_i c_i)_{i<\omega}$ for a global M-invariant type extending $\mathrm{tp}(bc/M)$, the b_i are distinct (like the columns of the TP$_2$-array). See also Proposition 6.3.20 below.

Definition 6.3.16. For subsets A, B, and C of $\mathbb{M}$, we define

$$A \mathop{\smile\!\!\!\!\!\!|}\limits^{M}_{C} B \iff \text{for all } C \subseteq C' \subseteq \operatorname{acl}(BC), A \mathop{\smile\!\!\!\!\!\!|}\limits^{a}_{C'} B.$$

We will show that $\mathop{\smile\!\!\!\!\!\!|}^M$ agrees with dividing independence $\mathop{\smile\!\!\!\!\!\!|}^d$ in $T_L^\emptyset$ (Proposition 6.3.17). The notation $\mathop{\smile\!\!\!\!\!\!|}^M$ comes from Adler [Adl05], who calls this relation "M-dividing independence". Adler shows that algebraic independence $\mathop{\smile\!\!\!\!\!\!|}^a$ satisfies all of his axioms for a strict independence relation except possibly base monotonicity (which it fails in $T_L^\emptyset$ whenever there is an n-ary function symbol, $n \geq 2$).

The relation $\mathop{\smile\!\!\!\!\!\!|}^M$ is obtained from $\mathop{\smile\!\!\!\!\!\!|}^a$ by forcing base monotonicity, and it satisfies all of the axioms of a strict independence relation except possibly local character and extension. If we go one step further and force extension, we get the relation $\mathop{\smile\!\!\!\!\!\!|}^þ$ of thorn forking independence [Adl05, Section 4], just as we get the relation $\mathop{\smile\!\!\!\!\!\!|}^f$ of forking independence by forcing extension on $\mathop{\smile\!\!\!\!\!\!|}^d$. But, as we will see, $\mathop{\smile\!\!\!\!\!\!|}^M$ already satisfies extension in $T_L^\emptyset$ (Proposition 6.3.18), so M-dividing independence, thorn forking independence, dividing independence, and forking independence all coincide in $T_L^\emptyset$. Of course, when L contains an n-ary function symbol with $n \geq 2$, these independence relations lack local character, since $T_L^\emptyset$ is not simple, so $T_L^\emptyset$ is also not rosy. In contrast, $\mathop{\smile\!\!\!\!\!\!|}^a = \mathop{\smile\!\!\!\!\!\!|}^K$ has local character but lacks base monotonicity. This tension between local character and base monotonicity is characteristic of the difference between forking independence and Kim-independence in NSOP_1 theories.

Proposition 6.3.17. *In $T_L^\emptyset$, $\mathop{\smile\!\!\!\!\!\!|}^d = \mathop{\smile\!\!\!\!\!\!|}^M$.*

Proof. In any theory, if $A \mathop{\smile\!\!\!\!\!\!|}^d_C B$, then $A \mathop{\smile\!\!\!\!\!\!|}^M_C B$ [Adl05, Remark 5.4.(4)]. So suppose $A \mathop{\smile\!\!\!\!\!\!|}^M_C B$. We may assume $B = \operatorname{acl}(BC) = \langle BC \rangle$, since $A \mathop{\smile\!\!\!\!\!\!|}^M_C B$ implies $A \mathop{\smile\!\!\!\!\!\!|}^M_C \operatorname{acl}(BC)$ and $A \mathop{\smile\!\!\!\!\!\!|}^d_C \operatorname{acl}(BC)$ implies $A \mathop{\smile\!\!\!\!\!\!|}^d_C B$.

Let b be a tuple enumerating B, and let $(b_i)_{i<\omega}$ be a C-indiscernible sequence, with $b_0 = b$. Let B_i be the set enumerated by b_i. Let C' be the set of all elements of B which appear in some b_i for $i \neq 0$. Then $C \subseteq C' \subseteq B$, every element of C' appears in every b_i, and $C' = \langle C' \rangle$. Letting c' enumerate C' and writing $b_i = (c', b_i')$ for all i, we have that $(b_i')_{i<\omega}$ is a C'-indiscernible sequence, and b_i' and b_j' are disjoint for all $i \neq j$.

Let the tuple x enumerate $\langle AC' \rangle \setminus C'$, and let the tuple y_0 enumerate $\langle AB \rangle \setminus (\langle AC' \rangle \cup B)$. By assumption, we have $A \mathop{\smile\!\!\!\!\!\!|}^a_{C'} B$, so (x, y_0, c', b_0') enumerates $\langle AB \rangle$ without repetitions (since no elements of x and b_0' are equal). Let $D = \langle \bigcup_{i<\omega} B_i \rangle$, and let d enumerate $D \setminus \bigcup_{i<\omega} B_i$.

Let $p(x, y_0, c', b_0') = \operatorname{diag}_f(\langle AB \rangle)$, and let $q(c', (b_i')_{i<\omega}, d) = \operatorname{diag}_f(D)$. Consider the flat diagram $q(c', (b_i')_{i<\omega}, d) \cup \bigcup_{i<\omega} p(x, y_i, c', b_i')$, where the y_i for $i > 0$ are new tuples. This is consistent, since any two copies of p agree on $\operatorname{diag}_f(\langle AC' \rangle)$, and the copy of p indexed by i agrees with q on $\operatorname{diag}_f(B_i)$. So by Lemma 6.3.4, there is an L-structure X with this diagram, and we can embed X in $\mathbb{M}$ over D by $i\colon X \to \mathbb{M}$, since $\operatorname{qftp}_X(c', (b_i')_{i<\omega}, d') = \operatorname{qftp}_{\mathbb{M}}(c', (b_i')_{i<\omega}, d')$.

Letting A' be the subset of $(i(x), c')$ corresponding to A as a subset of (x, c'), we have by quantifier elimination that $\mathrm{tp}(A'/B_i) = \mathrm{tp}(A/B)$ for all i. So $A \mathop{\smile\!\!\!\!\!\!\mid}^d_C B$. □

Proposition 6.3.18. *The relations* $\mathop{\smile\!\!\!\!\!\!\mid}^M$ *and* $\mathop{\smile\!\!\!\!\!\!\mid}^d$ *satisfy extension over arbitrary sets, so* $\mathop{\smile\!\!\!\!\!\!\mid}^M = \mathop{\smile\!\!\!\!\!\!\mid}^þ$ *and* $\mathop{\smile\!\!\!\!\!\!\mid}^d = \mathop{\smile\!\!\!\!\!\!\mid}^f$*.*

Proof. By Proposition 6.3.17, it suffices to show that $\mathop{\smile\!\!\!\!\!\!\mid}^M$ satisfies extension. Suppose we have $A \mathop{\smile\!\!\!\!\!\!\mid}^M_C B$, and let B' be another set. We may assume that $C \subseteq A$, $C \subseteq B \subseteq B'$, and A, C, B and B' are algebraically closed. We would like to show that there exists $A' \equiv_B A$ such that $A' \mathop{\smile\!\!\!\!\!\!\mid}^M_C B'$.

Let $D = \langle AB\rangle \otimes_B B'$, the fibered coproduct of $\langle AB\rangle$ and B' over B in the category of L-structures. We can give an explicit description of D: Let D_0 be the disjoint union of $\langle AB\rangle$ and B' over B, i.e. with the elements of B in $\langle AB\rangle$ and in B' identified. Any term with parameters from D_0 can be uniquely simplified with respect to $\langle AB\rangle$ and B', by iteratively replacing any function symbol whose arguments are all elements of $\langle AB\rangle$, or whose arguments are all elements of B', by its value (as usual, we view constant symbols as 0-ary functions). Then the underlying set of D is given by the simplified terms with parameters from D_0, i.e. those terms with the property that no function symbol appearing in the term has all its argument in $\langle AB\rangle$ or all its arguments in B'. The interpretation of function symbols in D is the obvious one (compose the function with the given simplified terms, then simplify if necessary), and the only instances of relations which hold in D are those which hold in $\langle AB\rangle$ or in B'.

Note that $\langle AB\rangle$ and B' embed in D, by sending an element a to the term $a \in D_0$, and $\langle AB\rangle \cap B' = B$ in D. Identifying these structures with their isomorphic copies in D, there is an embedding $i\colon D \to \mathbb{M}$ which is the identity on B'. Let $A' = i(A)$. Then $\langle A'B\rangle$ is isomorphic to $\langle AB\rangle$, so $A' \equiv_B A$. In particular, $A' \mathop{\smile\!\!\!\!\!\!\mid}^M_C B$. Of course, $\langle A'B'\rangle$ is isomorphic to $\langle A'B\rangle \otimes_B B'$.

Towards showing that $A' \mathop{\smile\!\!\!\!\!\!\mid}^M_C B'$, pick C' with $C \subseteq C' \subseteq B'$. We may assume that C' is algebraically closed. Let $\widetilde{C} = C' \cap B$, which is also algebraically closed. We will prove by induction on terms with parameters from $A'C'$ that:

1. If such a term evaluates to an element of B', then that element is in C'.
2. If such a term evaluates to an element of $\langle A'B\rangle$, then that element is in $\langle A'\widetilde{C}\rangle$.

First, we handle the base cases. The constant symbols are automatically in C' and in $\langle A'\widetilde{C}\rangle$. The parameters from C' are already in C', and if they are also in $\langle A'B\rangle$, then since $C' \subseteq B'$, they are in $\langle A'B\rangle \cap B' = B$, and hence in $\widetilde{C} \subseteq \langle A'\widetilde{C}\rangle$. On the other hand, the parameters from A' are already in $\langle A'\widetilde{C}\rangle$, and if they are also in B', then they are in $\langle A'B\rangle \cap B' = B$, hence in $A' \cap B = C$ (as $A' \mathop{\smile\!\!\!\!\!\!\mid}^a_C B$), and $C \subseteq C'$.

So suppose our term is $f(t_1, \ldots, t_n)$, where $t_1, \ldots, t_n$ are terms with parameters from $A'C'$ satisfying (1) and (2). Suppose t_i evaluates to c_i for all i, and let $b = f(c_1, \ldots, c_n)$.

Case 1: $c_i \in B'$ for all i. Then by induction, $c_i \in C'$ for all i, and hence $b \in C'$. And if b is also in $\langle A'B\rangle$, then by the argument for parameters in C', it is in $\langle A'\widetilde{C}\rangle$.

Case 2: $c_i \in \langle A'B\rangle$ for all i. Then by induction, $c_i \in \langle A'\widetilde{C}\rangle$ for all i. (2) is immediate, since also $b \in \langle A'\widetilde{C}\rangle$. For (1), suppose $b \in B'$. Then $b \in \langle A'B\rangle \cap B' = B$. But since $A' \mathop{\smile\!\!\!\!\!\!\mid}^M_C B$, we have $A' \mathop{\smile\!\!\!\!\!\!\mid}^a_{\widetilde{C}} B$, so $b \in \widetilde{C} \subseteq C'$.

Case 3: Neither of the above. Then writing each c_i in its normal form as a simplified term in $\langle A'B\rangle \otimes B'$, the element b does not simplify down to a single parameter from $\langle A'B\rangle$ or from B', since it is not the case that all of the arguments of f come from $\langle A'B\rangle$ or from B'. So (1) and (2) are trivially satisfied.

Of course, condition (1) establishes that $A' \mathop{\smile\!\!\!\!\!\!\mid}^a_{C'} B'$, as desired. □

Remark 6.3.19. If we define a new relation $\mathop{\smile\!\!\!\!\!\!\mid}^{\otimes}$ on subsets of $\mathbb{M}$ by $A \mathop{\smile\!\!\!\!\!\!\mid}^{\otimes}_C B$ if and only if the natural map $\langle AC\rangle \otimes_{\langle C\rangle} \langle BC\rangle \to \langle ABC\rangle$ is an isomorphism, then we can interpret the proof of Proposition 6.3.18 as a "mixed transitivity" statement. If $C \subseteq B \subseteq B'$, then:

$$A \mathop{\smile\!\!\!\!\!\!\mid}^{d}_{C} B \text{ and } A \mathop{\smile\!\!\!\!\!\!\mid}^{\otimes}_{B} B' \implies A \mathop{\smile\!\!\!\!\!\!\mid}^{d}_{C} B'.$$

Thanks to this form of transitivity, $\mathop{\smile\!\!\!\!\!\!\mid}^d$ inherits extension from $\mathop{\smile\!\!\!\!\!\!\mid}^{\otimes}$.

It may be worth noting that a similar pattern occurs in Conant's analysis of forking and dividing in the theory T_n of the generic K_n-free graph [Con14], which has SOP 3 (and hence also has SOP 1) when $n \geq 3$. Conant defines a relation $\mathop{\smile\!\!\!\!\!\!\mid}^R$ which satisfies extension ([Con14, Lemma 5.2]), and the proof of [Con14, Theorem 5.3] shows that $\mathop{\smile\!\!\!\!\!\!\mid}^d$ and $\mathop{\smile\!\!\!\!\!\!\mid}^R$ enjoy the same mixed transitivity property:

$$A \mathop{\smile\!\!\!\!\!\!\mid}^{d}_{C} B \text{ and } A \mathop{\smile\!\!\!\!\!\!\mid}^{R}_{B} B' \implies A \mathop{\smile\!\!\!\!\!\!\mid}^{d}_{C} B'.$$

As a consequence, $\mathop{\smile\!\!\!\!\!\!\mid}^d$ inherits extension from $\mathop{\smile\!\!\!\!\!\!\mid}^R$, and hence $\mathop{\smile\!\!\!\!\!\!\mid}^d = \mathop{\smile\!\!\!\!\!\!\mid}^f$ in T_n.

Proposition 6.3.18 tells us that forking equals dividing for complete types. On the other hand, forking does not equal dividing for formulas, even over models.

Proposition 6.3.20. *Suppose L contains an n-ary function symbol f with $n \geq 2$. For any set A, there is a formula which forks over A but does not divide over A.*

Proof. Let $b = (b_0, \ldots, b_{n-1})$ be an $(n-1)$-tuple such that $b_0 \notin \langle A\rangle$, and let c be an element such that $c \notin \langle Ab\rangle$. Then the formula $\varphi(x; b, c)$ given by $f(x, b) = c \vee f(x, x) = b_0$ forks over A but does not divide over A.

First, we show that the subformulas $f(x, b) = c$ and $f(x, x) = b_0$ divide over A. For the first, let $(b_i c_i)_{i<\omega}$ be any sequence of realizations of $\text{tp}(bc/A)$ such that $b_i = (b_{i,0}, \ldots, b_{i,n-1})$ and $b_i, j = b_j$ for all $i < \omega$ and $j < n$, but $c_i \neq c_j$ for all $i \neq j$ (this is possible, since $c \notin \text{acl}(Ab)$). Then $\{f(x, b) = c_i \mid i < \omega\}$ is 2-inconsistent. For the second, let $(b_i)_{i<\omega}$ be a sequence of realizations of $\text{tp}(b/A)$ with $b_i, 0 \neq b_j, 0$ for all $i \neq j$. Again, $\{f(x, x) = b_i \mid i < \omega\}$

is 2-inconsistent. Since $\varphi(x; b, c)$ is a disjunction of two formulas which divide over A, it forks over A.

To show that this formula does not divide, let $(b_i c_i)_{i<\omega}$ be any A-indiscernible sequence with $b_0 c_0 = bc$. If $b_i = b$ for all $i < \omega$, $\{\varphi(x; b_i, c_i) \mid i < \omega\}$ is consistent, witnessed by any a such that $f(a, a) = b$. On the other hand, if $b_i \neq b_j$ for all $i \neq j$, then we can find some a such that $f(a, b_i) = c_i$ for all $i < \omega$, so $\{\varphi(x; b_i, c_i) \mid i < \omega\}$ is consistent in this case too. □

Elimination of imaginaries

Definition 6.3.21. The theory T has *weak elimination of imaginaries* if for all imaginary elements e, there is a real element $a \in \operatorname{acl}^{\mathrm{eq}}(e)$ with $e \in \operatorname{dcl}^{\mathrm{eq}}(a)$.

We prove weak elimination of imaginaries for $T_L^{\emptyset}$. The argument follows the standard route to elimination of imaginaries via an independence theorem as in [Hru91, Proposition 3.1] and [CP98, Subsection 2.9].

Proposition 6.3.22. *$T_L^{\emptyset}$ has weak elimination of imaginaries.*

Proof. Suppose we are given an imaginary element e, and suppose a is a tuple from $\mathbb{M}$ and f is a 0-definable function (in $\mathbb{M}^{eq}$) with $f(a) = e$. Put $C = \operatorname{acl}^{\mathrm{eq}}(e) \cap \mathbb{M}$ and $q = \operatorname{tp}(a/C)$. We may assume $\operatorname{tp}(a/\operatorname{acl}^{eq}(e))$ is not algebraic, because, if it is, we're done.

Claim: There are $a, b \models q$ with $a \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} b$ with $f(a) = f(b) = e$.

Proof of claim: $\operatorname{tp}(a/\operatorname{acl}^{\mathrm{eq}}(e))$ is non-algebraic so, by extension, we can find $b \models \operatorname{tp}(a/\operatorname{acl}^{\mathrm{eq}}(e))$ with $b \mathop{\smile\!\!\!\!\!\vert}^{a}_{\operatorname{acl}^{\mathrm{eq}}(e)} a$ in $\mathbb{M}^{\mathrm{eq}}$. Note that also $f(b) = e$. Since $\operatorname{acl}^{\mathrm{eq}}(a) \cap \operatorname{acl}^{\mathrm{eq}}(b) = \operatorname{acl}^{\mathrm{eq}}(e)$, by intersecting with $\mathbb{M}$, we obtain $\operatorname{acl}(a) \cap \operatorname{acl}(b) = C$, that is, $a \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} b$. □

Let a, b be given as in the claim. If e is definable over C, we are done. If e is not definable over C, there is $e' \models \operatorname{tp}(e/C)$ with $e' \neq e$ and we can find $c', d' \models q$ with $c' \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} d'$ and $f(c') = f(d') = e'$, again by the claim. As $c' \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} d'$ we may, by extension, choose $c \equiv_{Cd'} c'$ with $c \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} ad'$. In particular, this gives $c \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} a$ and $f(c) \neq f(a)$.

As $a, c \models q$ we have $a \equiv_C c$. Moreover, we have $a \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} b$ and $c \mathop{\smile\!\!\!\!\!\vert}^{a}_{C} a$ so there is $a_* \models \operatorname{tp}(a/Cb) \cup \operatorname{tp}(c/Ca)$ by the independence theorem (Theorem 6.3.12). Then we have $f(a_*) = f(b) = f(a) \neq f(a_*)$, a contradiction. □

Remark 6.3.23. $T_L^{\emptyset}$ does not eliminate imaginaries, since it does not even code unordered pairs. That is, there is no definable binary function $f(x, y)$ such that $f(a, b) = f(c, d)$ if and only if $\{a, b\} = \{c, d\}$. To see this, note that, by quantifier elimination, every definable function is defined piecewise by terms. Let F_2 be the L-structure freely generated over $\langle\emptyset\rangle$ (the substructure generated by the constants) by two elements, a and b. Then if t is a term, considered in the variable context $\{x, y\}$, such that $t(a, b) = t(b, a)$, then t does not mention the variables, i.e. t evaluates to an element of $\langle\emptyset\rangle$. For any copy of F_2 embedded in $\mathbb{M}$, $\operatorname{tp}(a, b) = \operatorname{tp}(b, a)$, since the automorphism of F_2 swapping a and b extends to an automorphism of $\mathbb{M}$. So any function f coding unordered pairs must be defined by the same term t on (a, b) and on (b, a). Then $t(a, b) = t(b, a) \in \langle\emptyset\rangle$. But this is a contradiction, since there are automorphisms of $\mathbb{M}$ which do not fix $\{a, b\}$ setwise.

6.4 Generic expansion and Skolemization

The theories T_{Sk} and $T_{L'}$

Definition 6.4.1. Given a language L, define the language L_{Sk} by adding to L, for each formula $\varphi(x;y)$ with $l(x) = 1$, an $l(y)$-ary function symbol f_φ. The *Skolem expansion of T* is the L_{Sk}-theory T_+ defined by

$$T_+ = T \cup \{\forall y\, (\exists x\, \varphi(x;y) \to \varphi(f_\varphi(y);y)) \mid \varphi(x;y) \in L, l(x) = 1\}.$$

Note that the Skolem expansion of T contains Skolem functions for every formula of L, but does *not* contain Skolem functions for every formula of L_{Sk}.

Fact 6.4.2. Suppose T is a model complete theory in the language L that eliminates the quantifier $\exists^\infty$.

1. The Skolem expansion T_+ of T has a model completion T_{Sk} [Win75, Theorem 2]. We will refer to the theory T_{Sk} as the *generic Skolemization of T*.
2. For any language L' containing L, T, considered as an L'-theory, has a model companion $T_{L'}$ [Win75, Theorem 5]. We will refer to the theory $T_{L'}$ as the *generic L'-expansion of T*.

Fact 6.4.3. [Win75, Corollary 3 to Theorem 4] Under the hypotheses of Fact 6.4.2, T_{Sk} also eliminates the quantifier $\exists^\infty$.

Preservation of NSOP_1

In this subsection, suppose T is a fixed model complete theory in the language L that eliminates the quantifier $\exists^\infty$, and let L' be an arbitrary language containing L. We may choose monster models $\mathbb{M}_{\mathrm{Sk}} \models T_{Sk}$ and $\mathbb{M}_{L'} \models T_{L'}$. We may assume that both monster models have a common reduct to a monster model $\mathbb{M} \models T$. Note that we do not assume T is complete, so reasoning in $\mathbb{M}$ amounts to working in an arbitrary completion of T.

Definition 6.4.4. Let $\mathbb{M}^*$ denote either $\mathbb{M}_{\mathrm{Sk}}$ or $\mathbb{M}_{L'}$ and let L^* denote the corresponding language, either L_{Sk} or L'. For $a, b \in \mathbb{M}^*$ and $M \prec_{L^*} \mathbb{M}^*$, define

$$a \mathop{\downarrow}^{*}_{M} b \iff \mathrm{acl}_{L^*}(Ma) \mathop{\downarrow}^{K}_{M} \mathrm{acl}_{L^*}(Mb) \text{ in } \mathbb{M}.$$

Theorem 6.4.5. *Let $\mathbb{M}^*$ denote either $\mathbb{M}_{Sk}$ or $\mathbb{M}_{L'}$ and let L^* denote the corresponding language, either L_{Sk} or L'. If T is NSOP_1, then $\mathop{\downarrow}^*$ satisfies the independence theorem.*

Proof. We're given $M \prec \mathbb{M}^*$ and tuples a, a', b, c, with $a \mathop{\downarrow}^*_M b$, $a' \mathop{\downarrow}^*_M c$, $b \mathop{\downarrow}^*_M c$, and $a \equiv_M a'$. Let x_M be a tuple enumerating M, let x_a, $x_{a'}$, x_b and x_c be tuples enumerating $\mathrm{acl}_{L^*}(Ma) \setminus M$, $\mathrm{acl}_{L^*}(Ma') \setminus M$, $\mathrm{acl}_{L^*}(Mb) \setminus M$, and $\mathrm{acl}_{L^*}(Mc) \setminus M$, respectively, and let x_{ab}, $x_{a'c}$, and x_{bc}

be tuples enumerating $\mathrm{acl}_{L^*}(Mab) \setminus Mx_ax_b$, $\mathrm{acl}_{L^*}(Ma'c) \setminus Mx_{a'}x_c$, and $\mathrm{acl}_{L^*}(Mbc) \setminus Mx_bx_c$, respectively.

We have $x_a \mathop{\smile\!\!\!\!\!\mid}^K_M x_b$, $x_{a'} \mathop{\smile\!\!\!\!\!\mid}^K_M x_c$, and $x_b \mathop{\smile\!\!\!\!\!\mid}^K_M x_c$ in $\mathbb{M}$. By the algebraically reasonable independence theorem in T, we can find $x_{a''}$ in $\mathbb{M}$ such that $x_{a''} \equiv^L_{Mx_b} x_a$, $x_{a''} \equiv^L_{Mx_c} x_{a'}$, $x_{a''} \mathop{\smile\!\!\!\!\!\mid}^K_M x_bx_c$, and further $x_{a''} \mathop{\smile\!\!\!\!\!\mid}^a_{Mx_b} x_c$, $x_{a''} \mathop{\smile\!\!\!\!\!\mid}^a_{Mx_c} x_b$, and $x_b \mathop{\smile\!\!\!\!\!\mid}^a_{Mx_{a''}} x_c$. By algebraically reasonable extension, at the expense of moving $x_{a''}$ by an automorphism in $\mathrm{Aut}(\mathbb{M}/Mx_bx_c)$, we may assume that $x_{a''} \mathop{\smile\!\!\!\!\!\mid}^K_M x_bx_cx_{bc}$ and $x_{a''} \mathop{\smile\!\!\!\!\!\mid}^a_{Mx_bx_c} x_{bc}$.

Pick an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}/Mx_b)$ moving x_a to $x_{a''}$, and set $x_{a''b} = \sigma(x_{ab})$, so $x_{a''}x_{a''b} \equiv^L_{Mx_b} x_ax_{ab}$. Similarly, pick an automorphism $\sigma' \in \mathrm{Aut}(\mathbb{M}/Mx_c)$ moving $x_{a'}$ to $x_{a''}$, and set $x_{a''c} = \sigma'(x_{a'c})$, so $x_{a''}x_{a''c} \equiv^L_{Mx_c} x_{a'}x_{a'c}$. Now there are subtuples $y_{a''b} \subseteq x_{a''b}$ and $y_{a''c} \subseteq x_{a''c}$ enumerating $\mathrm{acl}_L(Mx_{a''}x_b) \setminus Mx_{a''}x_b$ and $\mathrm{acl}_L(Mx_{a''}x_c) \setminus Mx_{a''}x_c$, respectively. The algebraic independencies obtained so far imply that the tuples x_M, $x_{a''}$, x_b, x_c, $y_{a''b}$, $y_{a''c}$, and x_{bc} are pairwise disjoint.

By two applications of extension for algebraic independence over algebraically closed bases, we can find tuples $z_{a''b}$ and $z_{a''c}$ such that $z_{a''b} \equiv^L_{Mx_{a''}x_by_{a''b}} x_{a''b}$ and $z_{a''c} \equiv^L_{Mx_{a''}x_cy_{a''c}} x_{a''c}$, and so that the tuples x_M, $x_{a''}$, x_b, x_c, $z_{a''b}$, $z_{a''c}$, and x_{bc} are pairwise disjoint.

Let $\widehat{M} \prec \mathbb{M}$ be a small model of T containing all these tuples. We will expand $\widehat{M}$ to an L^*-structure, in order to embed it in $\mathbb{M}^*$. Let $p_{a''b} = \mathrm{diag}_f(\mathrm{acl}_{L^*}(Mab))$, $p_{a''c} = \mathrm{diag}_f(\mathrm{acl}_{L^*}(Ma'c))$, and $p_{bc} = \mathrm{diag}_f(\mathrm{acl}_{L^*}(Mbc))$. We define interpretations of the relations, functions, and constants of L^* according to

$$p_{a''b}(x_M, x_{a''}, x_b, z_{a''b}) \cup p_{a''c}(x_M, x_{a''}, x_c, z_{a''c}) \cup p_{bc}(x_M, x_b, x_c, x_{bc}),$$

for the tuples that these diagrams refer to. This is consistent, since $p_{a''b}$, $p_{a''c}$, and p_{bc} agree on $\mathrm{diag}_f(\mathrm{acl}_{L^*}(Ma)) = \mathrm{diag}_f(\mathrm{acl}_{L^*}(Ma'))$ (allowing $x_{a''}$ to enumerate both $\mathrm{acl}_{L^*}(Ma) \setminus M$ and $\mathrm{acl}_{L^*}(Ma') \setminus M$), $\mathrm{diag}_f(\mathrm{acl}_{L^*}(Mb))$, and $\mathrm{diag}_f(\mathrm{acl}_{L^*}(Mc))$. In the case that $L^* = L_{\mathrm{Sk}}$, we observe that the values of the functions specified by these diagrams really give Skolem functions, since we have preserved the underlying L-types of all the tuples. For tuples not referred to by these diagrams, we define the interpretations of the relations and functions arbitrarily, taking care in the case that $L^* = L_{Sk}$ to satisfy the Skolem axioms (this is always possible, since $\widehat{M}$ is a model).

Having expanded $\widehat{M}$ to an L^*-structure, we can embed it in $\mathbb{M}^*$ by $i\colon \widehat{M} \to \mathbb{M}^*$. Further, we may assume that i is the identity on (x_M, x_b, x_c, x_{bc}), since

$$\mathrm{qftp}^{L^*}_{\widehat{M}}(x_M, x_b, x_c, x_{bc}) = \mathrm{qftp}^{L^*}_{\mathbb{M}^*}(x_M, x_b, x_c, x_{bc}).$$

Let a'' be the subtuple of $i(x_M, x_{a''})$ corresponding to the subtuple of (x_M, x_a) enumerating a.

Now $\mathrm{qftp}(x_M, i(x_{a''}), x_b, i(z_{a''b})) = \mathrm{qftp}(x_M, x_a, x_b, x_{ab})$, so $a'' \equiv_{Mb} a$, and similarly

$$\mathrm{qftp}(x_M, i(x_{a''}), x_c, i(z_{a''c})) = \mathrm{qftp}(x_M, x_{a'}, x_c, x_{a'c}),$$

so $a'' \equiv_{Mc} a'$. And $a'' \mathop{\smile\!\!\!\!\!\mid}^*_M bc$, since $i(x_{a''})$ (which enumerates $\mathrm{acl}_{L^*}(Ma'') \setminus M$) is Kim-independent over M from $x_bx_cx_{bc}$ (which enumerates $\mathrm{acl}_{L^*}(Mbc) \setminus M$). $\square$

Corollary 6.4.6. *Suppose T is* NSOP$_1$ *Then:*

1. *T_{Sk} is* NSOP$_1$*, and $\mathop{\smile\!\!\!\!\!\!\mid}\,^K = \mathop{\smile\!\!\!\!\!\!\mid}\,^*$.*

2. *$T_{L'}$ is* NSOP$_1$*, and $\mathop{\smile\!\!\!\!\!\!\mid}\,^K = \mathop{\smile\!\!\!\!\!\!\mid}\,^*$.*

Proof. We use Theorem 6.2.4. By Theorem 6.4.5, $\mathop{\smile\!\!\!\!\!\!\mid}\,^*$ satisfies the independence theorem. Existence over models, monotonicity, and symmetry follow immediately from the definition and the corresponding properties of $\mathop{\smile\!\!\!\!\!\!\mid}\,^K$ in the reduct.

Now suppose $a \not\mathop{\smile\!\!\!\!\!\!\mid}[*]_M b$, witnessed by $a' \in \mathrm{acl}_{L^*}(Ma)$ and $b' \in \mathrm{acl}_{L^*}(Mb)$ such that $a' \not\mathop{\smile\!\!\!\!\!\!\mid}[K]_M b'$ in the reduct. Let $\varphi(x'; b', m)$ be the L-formula given strong finite character for $\mathop{\smile\!\!\!\!\!\!\mid}\,^K$ in the reduct. Let $\chi(x', a, m_1)$ isolate $\mathrm{tp}_{L^*}(a'/Ma)$, and let $\psi(y', b, m_2)$ isolate $\mathrm{tp}_{L^*}(b'/Mb)$. Then the formula $\varphi'(x; b, m, m_1, m_2)$ given by $\exists x' \, \exists y' \, (\chi(x', x, m_1) \wedge \psi(y', b, m_2) \wedge \varphi(x', y', m))$ gives strong finite character for $\mathop{\smile\!\!\!\!\!\!\mid}\,^*$. Similarly, if φ gives witnessing for $\mathop{\smile\!\!\!\!\!\!\mid}\,^K$ in the reduct, the same formula φ' gives witnessing for $\mathop{\smile\!\!\!\!\!\!\mid}\,^*$. $\square$

Remark 6.4.7. By [Nüb04, Lemma 3.1], T_{Sk} will never be a simple theory.

Iterating to get built-in Skolem functions

Definition 6.4.8. An L-theory T has built-in Skolem functions if, for every formula $\varphi(x; y) \in L$ with $l(x) = 1$, there is a definable $l(y)$-ary function f_φ so that

$$T \models \forall y \, (\exists x \, \varphi(x; y) \to \varphi(f_\varphi(y); y)).$$

Corollary 6.4.9. *Any* NSOP$_1$ *theory T which eliminates $\exists^\infty$ has an expansion to an* NSOP$_1$ *theory T^∞_{Sk} in a language L^∞_{Sk} with built-in Skolem functions. Moreover, if $\mathbb{M}^\infty_{\mathrm{Sk}}$ is a monster model for T^∞_{Sk}, and $\mathbb{M}$ is its reduct to L, then for every $M \prec \mathbb{M}^\infty_{\mathrm{Sk}}$ and tuples a and b,*

$$a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b \text{ in } \mathbb{M}^\infty_{\mathrm{Sk}} \iff \mathrm{acl}_{L^\infty_{\mathrm{Sk}}}(Ma) \mathop{\smile\!\!\!\!\!\!\mid}^K_M \mathrm{acl}_{L^\infty_{\mathrm{Sk}}}(Mb) \text{ in } \mathbb{M}.$$

Proof. We define T^∞_{Sk} by induction. Let T_0 be the Morleyization of T in the expanded language L_0, so T_0 is model complete. Given T_n, which we may assume by induction to be a model complete NSOP$_1$ theory which eliminates $\exists^\infty$, let $L_{n+1} = (L_n)_{\mathrm{Sk}}$, and let $T_{n+1} = (T_n)_{\mathrm{Sk}}$. Then, by Fact 6.4.3 and Theorem 6.4.5, T_{n+1} is again a model complete NSOP$_1$ theory which eliminates $\exists^\infty$ and has Skolem function for formulas in the language of T_n. And by Theorem 6.4.5 and induction, if $\mathbb{M}_{n+1}$ is a monster model for T_{n+1}, and $\mathbb{M}$ is its reduct to L,

$$\begin{aligned} a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b \text{ in } \mathbb{M}_{n+1} &\iff \mathrm{acl}_{L_{n+1}}(Ma) \mathop{\smile\!\!\!\!\!\!\mid}^K_M \mathrm{acl}_{L_{n+1}}(Mb) \text{ in } \mathbb{M}_n \\ &\iff \mathrm{acl}_{L_{n+1}}(Ma) \mathop{\smile\!\!\!\!\!\!\mid}^K_M \mathrm{acl}_{L_{n+1}}(Mb) \text{ in } \mathbb{M}, \end{aligned}$$

since a set which is $\mathrm{acl}_{L_{n+1}}$-closed is also acl_{L_n}-closed.

Then the theory $T^{\infty}_{\mathrm{Sk}} = \bigcup_{n<\omega} T_n$ is NSOP$_1$ and has built-in Skolem functions. It remains to lift the characterization of Kim-independence to this theory.

If $\varphi(x;b)$ is a formula in L^{∞}_{Sk}, possibly with parameters in M, there is some n so that $\varphi(x;y) \in L_n$. Let $I = (b_i)_{<\omega}$ be an M-finitely satisfiable sequence in $\mathrm{tp}_{L^{\infty}_{\mathrm{Sk}}}(b/M)$. Then I is also an M-finitely satisfiable sequence in $\mathrm{tp}_{L_n}(b/M)$. By Kim's lemma for Kim-dividing Theorem III.3.15, $\varphi(x;b)$ Kim-divides over M if and only if $\{\varphi(x;b_i) \mid i < \omega\}$ is consistent, hence $\varphi(x;b)$ Kim-divides over M in $\mathbb{M}_n$ if and only if $\varphi(x;b)$ Kim-divides over M in $\mathbb{M}^{\infty}_{\mathrm{Sk}}$. It follows that $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ in $\mathbb{M}^{\infty}_{\mathrm{Sk}}$ if and only if $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ in $\mathbb{M}_n$ for all n.

Now we show $a \mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ in $\mathbb{M}^{\infty}_{\mathrm{Sk}}$ if and only if $\mathrm{acl}_{L^{\infty}_{\mathrm{Sk}}}(Ma) \mathop{\smile\!\!\!\!\!\!\mid}^K_M \mathrm{acl}_{L^{\infty}_{\mathrm{Sk}}}(Mb)$ in $\mathbb{M}$. If

$$\mathrm{acl}_{L^{\infty}_{\mathrm{Sk}}}(Ma) \underset{M}{\overset{K}{\not\mathop{\smile\!\!\!\!\!\!\mid}}} \mathrm{acl}_{L^{\infty}_{\mathrm{Sk}}}(Mb)$$

in $\mathbb{M}$, then there is n so that $\mathrm{acl}_{L_n}(Ma) \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M \mathrm{acl}_{L_n}(Mb)$ hence there is some formula $\varphi(x;b) \in \mathrm{tp}_{L_n}(a/Mb)$ that Kim-divides over M in $\mathbb{M}_n$. This formula witnesses $a \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ in $\mathbb{M}^{\infty}_{\mathrm{Sk}}$. Conversely, if $a \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ in $\mathbb{M}^{\infty}_{\mathrm{Sk}}$, then there is some formula $\varphi(x;b) \in \mathrm{tp}_{L^{\infty}_{\mathrm{Sk}}}(a/Mb)$ witnessing this. Then for some n, $\varphi(x;b) \in \mathrm{tp}_{L_n}(a/Mb)$ and this formula witnesses $a \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M b$ in $\mathbb{M}_n$. It follows that $\mathrm{acl}_{L_n}(Ma) \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M \mathrm{acl}_{L_n}(Mb)$ in $\mathbb{M}$, and therefore $\mathrm{acl}_{L^{\infty}_{\mathrm{Sk}}}(Ma) \not\mathop{\smile\!\!\!\!\!\!\mid}^K_M \mathrm{acl}_{L^{\infty}_{\mathrm{Sk}}}(Mb)$ by monotonicity. $\square$

Chapter 7

SOP_1 and SOP_2 in PAC fields

Theories of fields have provided a fertile testing ground for many notions in abstract model theory and, among fields, the pseudo-algebraically closed fields were especially significant examples for simplicity theory and beyond. A field K is pseudo-algebraically closed (PAC) if every absolutely irreducible variety defined over K has a K-rational point. This class of fields was isolated by Ax in his celebrated characterization of pseudo-finite fields. Ax showed that an infinite field K is a model of the theory of finite fields if and only if K is perfect, the absolute Galois group of K is $\hat{\mathbb{Z}}$, and K is PAC. Cherlin, van den Dries, and Macintyre gave a comprehensive study of PAC fields, characterizing the completions of the completions of the theory of PAC fields and giving a description of types [CvdDM80]. The central observation underlying their work was that definable sets are controlled by algebraic data, encoded by sets definable in the stable separable or algebraic closure, together with first-order properties of the inverse system of the absolute Galois group.

Duret had observed that a PAC field which is not separably closed has the independence property and therefore is not stable, and a graph-coding construction of Cherlin, van Den Dries, and Macintyre had showed that theories of PAC fields can, in general, be as wild as possible. In subsequent developments, much attention was given to the problem of finding meaningful descriptions of the tame theories of PAC fields and connections were found between classification-theoretic properties of a PAC field and properties of its absolute Galois group. For example, one of the most satisfying results in this vein is the characterization of simple PAC fields. A field is called *bounded* if it has finitely many extensions of degree n for all n. Hrushovski showed that a perfect bounded PAC field has simple (in fact, supersimple) theory [Hru91] and later Chatzidakis and Pillay extended this result, establishing the simplicity of possibly imperfect bounded PAC fields [CP98]. Chatzidakis later proved that an unbounded PAC field has the tree property of the second kind, hence a PAC field is simple if and only if it is bounded [Cha99].

A remarkable theorem of Chatzidakis provides a powerful engine for proving theorems of this kind. She shows that, in a precise sense, independent amalgamation in a PAC field is completely controlled by independent amalgamation in the Galois group [Cha17]. Many of the important classification-theoretic dividing lines between simplicity and NSOP can be understood as describing the constraints on amalgamation within a given theory and, she shows, for $n \geq 3$, that if the theory of the inverse system of the absolute Galois group of a PAC field is NSOP_n then so is the field. The SOP_n hierarchy as initially defined by Shelah begins at $n = 3$ but Džamonja and Shelah later introduced SOP_1 and SOP_2, which are defined in a very different way from the rest of the hierarchy. Consequently, a new argument was needed to extend Chatzidakis's result to handle SOP_1 and SOP_2 (also known as TP_1). In this short chapter, we explain how to use the results of this thesis to handle these cases.

To show that a PAC field with NSOP_1 Galois group (in the inverse system language) has NSOP_1 theory, we show that Kim-independence in the field can be characterized in terms of Kim-independence in the Galois group. In the NSOP_2 case, a theory of independence is not available but we can instead make use of strongly indiscernible trees to reduce the analysis of possible witnesses to SOP_2 to concrete amalgamation problems, which in turn reduce to amalgamation problems in the Galois group. In both cases, the work we do here amounts to reformulating the problem of showing a field is NSOP_n to one that Chatzidakis's amalgamation theorem shows how to solve.

7.1 Preliminaries

Model theory

The properties we will be interested in are SOP_1 and SOP_2:

Definition 7.1.1. A formula $\varphi(x; y)$ has SOP_1 if there is a collection of tuples $(a_\eta)_{\eta \in 2^{<\omega}}$ satisfying the following:

- For all $\eta \in 2^\omega$, $\{\varphi(x; a_{\eta|n}) : n < \omega\}$ is consistent.
- For all $\eta, \nu \in 2^{<\omega}$, if $(\eta \wedge \nu) \frown \langle 0 \rangle \trianglelefteq \eta$ and $(\eta \wedge \nu) \frown \langle 1 \rangle = \nu$, then $\{\varphi(x; a_\eta), \varphi(x; a_\nu)\}$ is inconsistent.

A theory T has SOP_1 if some formula has SOP_1 modulo T. We say T is NSOP_1 if does not have SOP_1.

Definition 7.1.2. A formula $\varphi(x; y)$ has SOP_2 if there is a collection of tuples $(a_\eta)_{\eta \in 2^{<\omega}}$ satisfying the following:

- For all $\eta \in 2^\omega$, $\{\varphi(x; a_{\eta|n}) : n < \omega\}$ is consistent.
- For all $\eta \perp \nu \in 2^{<\omega}$, $\{\varphi(x; a_\eta), \varphi(x; a_\nu)\}$ is inconsistent.

A theory T has SOP_2 if some formula has SOP_2 modulo T. We say T is NSOP_2 if does not have SOP_2.

We will also make use of generalized indiscernibles. For a detailed treatment, see Chapter 1.

Definition 7.1.3. Suppose I is an L'-structure, where L' is some language.

1. We say $(a_i : i \in I)$ is a set of *I-indexed indiscernibles* if whenever $(s_0, \ldots, s_{n-1})$, $(t_0, \ldots, t_{n-1})$ are tuples from I with
$$\mathrm{qftp}_{L'}(s_0, \ldots, s_{n-1}) = \mathrm{qftp}_{L'}(t_0, \ldots, t_{n-1}),$$
then we have
$$\mathrm{tp}(a_{s_0}, \ldots, a_{s_{n-1}}) = \mathrm{tp}(a_{t_0}, \ldots, a_{t_{n-1}}).$$

2. In the case that $L' = \langle \trianglelefteq, <_{lex}, \wedge \rangle$ and $I = \kappa^{<\lambda}$, viewed as an L'-structure in the natural way, then an I-indexed indiscernible is called a *strongly indiscernible tree*.

3. We say that I-indexed indiscernibles have the *modeling property* if, given any $(a_i : i \in I)$ from $\mathbb{M}$, there is an I-indexed indiscernible $(b_i : i \in I)$ in $\mathbb{M}$ *locally based* on $(a_i : i \in I)$ – i.e., given any finite set of formulas Δ from L and a finite tuple $(t_0, \ldots, t_{n-1})$ from I, there is a tuple $(s_0, \ldots, s_{n-1})$ from I so that
$$\mathrm{qftp}_{L'}(t_0, \ldots, t_{n-1}) = \mathrm{qftp}_{L'}(s_0, \ldots, s_{n-1})$$
and also
$$\mathrm{tp}_\Delta(b_{t_0}, \ldots, b_{t_{n-1}}) = \mathrm{tp}_\Delta(a_{s_0}, \ldots, a_{s_{n-1}}).$$

The following was shown in Chapter 1:

Fact 7.1.4. If $\varphi(x, y)$ has SOP$_2$, then there is a strongly indiscernible tree $(a_\eta)_{\eta \in 2^{<\omega}}$ witnessing this.

Fields

Definition 7.1.5. Suppose F is a field. We say F is *pseudo-algebraically closed* (PAC) if every absolutely irreducible variety over F has an F-rational point.

Definition 7.1.6. Suppose G is a profinite group. Let $\mathcal{N}(G)$ be the collection of open normal subgroups of G. We define
$$\mathcal{S}(G) = \coprod_{N \in \mathcal{N}(G)} G/N.$$

Let L_G the language with a sort X_n for each $n \in \mathbb{Z}^+$, two binary relation symbols $\leq$, C, and a ternary relation P. We regard $\mathcal{S}(G)$ as an L_G-structure in the following way:

- The coset gN is in sort X_n if and only if $[G : N] \leq n$.
- $gN \leq hM$ if and only if $N \subseteq M$

- $C(gN, hM) \iff N \subseteq M$ and $gM = hM$.
- $P(g_1N_1, g_2N_2, g_3N_3) \iff N_1 = N_2 = N_3$ and $g_1g_2N_1 = g_3N_1$.

Note that we do not require that the sorts be disjoint (see [Cha98, Section 1] for a discussion on the syntax of this structure).

If K is a field, we write K^s for the separable closure of K, K^{alg} for the algebraic closure of K, and we write $\mathcal{G}(K)$ for the absolute Galois group of K. Interpretability of $\mathcal{S}(\mathcal{G}(F))$ in (F^{alg}, F) is proved in [Cha02, Proposition 5.5]. The "moreover" clause is clear from the proof.

Fact 7.1.7. Both F and $\mathcal{S}(\mathcal{G}(F))$ are interpretable in (K, F) where K is any algebraically closed field containing F. Call the interpretation π. Moreover, if $L \subseteq F$ is a subfield so that F is a regular extension of L, then the restriction of π to (K, L) produces an interpretation of $\mathcal{S}(\mathcal{G}(L))$, contained in $\mathcal{S}(\mathcal{G}(F))$ in a natural way.

Lemma 7.1.8. *Let F be a large sufficiently saturated and homogeneous field (i.e. a monster model of its theory) and $M \prec F$ a small elementary substructure. Suppose $A = acl(A)$, $B = acl(B)$ are subsets of F with $M \subseteq A \cap B$.*

1. *If $A \equiv_M B$ in F, then $\mathcal{S}(\mathcal{G}(A)) \equiv_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(B))$.*
2. *If $(A_i)_{i \in I}$ is an I-indexed indiscernible over M in $tp(A/M)$, then $(\mathcal{S}(\mathcal{G}(A_i)))_{i \in I}$ is an I-indexed indiscernible over $\mathcal{S}(\mathcal{G}(M))$.*
3. *If $A \mathop{\smile\!\!\!\!\!\!|}^u_M B$ in F, then $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!|}^u_{\mathcal{S}(\mathcal{G}(B))} \mathcal{S}(\mathcal{G}(B))$ in $\mathcal{S}(\mathcal{G}(M))$.*

Proof. (1) If $A \equiv_M B$ in F, then there is an automorphism $\sigma \in \operatorname{Aut}(F/M)$ with $\sigma(A) = B$. The map σ has an extension $\tilde{\sigma}$ to F^{alg} which is, then, an automorphism of the pair (F^{alg}, F) taking A to B and fixing M pointwise. It follows $A \equiv_M B$ in the pair (F^{alg}, F). Since $A = \operatorname{acl}(A)$ and $B = \operatorname{acl}(B)$, we know F is a regular extension of A and of B (see, e.g., [Cha99, Section 1.17]). By Fact 7.1.7, we have $\mathcal{S}(\mathcal{G}(A)) \equiv_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(B))$.

(2) Suppose $(A_i)_{i \in I}$ is an I-indexed indiscernible over M, given two k-tuples $\bar{i} = (i_0, \ldots, i_{k-1})$ and $\bar{j} = (j_0, \ldots, j_{k-1})$ from I with $\operatorname{qftp}(\bar{i}) = \operatorname{qftp}(\bar{j})$, we know $A_{i_0} \ldots A_{i_{k-1}} \equiv_M A_{j_0} \ldots A_{j_{k-1}}$ so $\operatorname{acl}(A_{i_0} \ldots A_{i_{k-1}}) \equiv_M \operatorname{acl}(A_{j_0} \ldots A_{j_{k-1}})$. Then by (1) $\mathcal{S}(\mathcal{G}(\operatorname{acl}(A_{i_0} \ldots A_{i_{k-1}}))) \equiv_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(\operatorname{acl}(A_{j_0} \ldots A_{j_{k-1}})))$, which implies $(\mathcal{S}(\mathcal{G}(A_i)))_{i \in I}$ is an I-indexed indiscernible over $\mathcal{S}(\mathcal{G}(M))$.

(3) In *any* theory, if π is an interpretation of the structure X in the structure Y, and $A \mathop{\smile\!\!\!\!\!\!|}^u_C B$ in Y, then $\pi(A) \mathop{\smile\!\!\!\!\!\!|}^u_{\pi(C)} \pi(B)$. It follows that if $A \mathop{\smile\!\!\!\!\!\!|}^u_M B$ in F, then $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\!\!\!\!\!\!|}^u_{\mathcal{S}(\mathcal{G}(M))} \mathcal{S}(\mathcal{G}(\mathcal{B}))$ by Fact 7.1.7. $\square$

A theorem of Chatzidakis

Theorem 7.1.9. *[Cha17] Let F be a PAC field, $E \prec F$, and A, B, C_1, and C_2 be algebraically closed subsets of F, with E contained in A, B, C_1, and C_2. Assume that $A \cap B = E$, that $A \mathop{\smile\!\!\!\!\!\!\,|}_E^{SCF} C_1$, $B \mathop{\smile\!\!\!\!\!\!\,|}_E^{SCF} C_2$. Assume further that there is an E^s-isomorphism $\varphi : C_1^s \to C_2^s$ such that $\varphi(C_1) = C_2$, and that there is $S_0 \subseteq S(\mathcal{G}(F))$ and L_G-elementary isomorphisms*

$$\begin{aligned} S\Psi_1 : \langle S(\mathcal{G}(C_1)), S(\mathcal{G}(A))\rangle &\to \langle S_0, S(\mathcal{G}(A))\rangle \\ S\Psi_2 : \langle S(\mathcal{G}(C_2)), S(\mathcal{G}(B))\rangle &\to \langle S_0, S(\mathcal{G}(B))\rangle. \end{aligned}$$

such that

1. *$S\Psi_1$ is the identity on $S(\mathcal{G}(A))$, $S\Psi_2$ is the identity on $S(\mathcal{G}(B))$, $S\Psi_i(S(\mathcal{G}(C_i))) = S_0$ for $i = 1, 2$,*

2. *If $S\Phi : S(\mathcal{G}(C_1)) \to S(\mathcal{G}(C_2))$ is the morphism double dual to φ, then*

$$S\Psi_2 S\Phi = S\Psi_1|_{S(\mathcal{G}(C_1))}.$$

Then, in some elementary extension F^* of F, there is C with $S(\mathcal{G}(C)) = S_0$ such that $C \equiv_A C_1$, $C \equiv_B C_2$ and $C \mathop{\smile\!\!\!\!\!\!\,|}_E^{SCF} AB$.

What we will use in the following arguments is actually a consequence of this theorem, which may easily be deduced from it: suppose F is a sufficiently saturated PAC field and $E \prec F$ is an elementary submodel. Suppose further that A, B, C_1, and C_2 be algebraically closed subsets of F, with E contained in A, B, C_1, and C_2. Assume that $A \cap B = E$, that $A \mathop{\smile\!\!\!\!\!\!\,|}_E^{\mathrm{SCF}} C_1$, $B \mathop{\smile\!\!\!\!\!\!\,|}_E^{\mathrm{SCF}} C_2$, and $C_1 \equiv_E C_2$. Then if in $\mathrm{Th}(S(\mathcal{G}(F)))$, there is S_0 with $S_0 \equiv_{S(\mathcal{G}(A))} S(\mathcal{G}(C_1))$ and $S_0 \equiv_{S(\mathcal{G}(B))} S(\mathcal{G}(C_2)))$, then there is C with $S(\mathcal{G}(C)) = S_0$ such that $C \equiv_A C_1$, $C \equiv_B C_2$ and $C \mathop{\smile\!\!\!\!\!\!\,|}_E^{\mathrm{SCF}} AB$.

7.2 SOP$_1$ and SOP$_2$

SOP$_1$

To begin, we will recall the definitions and main facts for Kim-independence. These are given a detailed treatment in Chapter 3.

Definition 7.2.1. Suppose $q(y)$ is an A-invariant global type. The formula $\varphi(x; y)$ q-divides over A if for some (equivalently, any) Morley sequence $\langle b_i : i < \omega\rangle$ in q over A, $\{\varphi(x; b_i) : i < \omega\}$ is inconsistent.

Definition 7.2.2. We say that a formula $\varphi(x; b)$ *Kim-divides* over A if there is some A-invariant global type $q \supseteq \mathrm{tp}(b/A)$ so that $\varphi(x; y)$ q-divides. The formula $\varphi(x; b)$ *Kim-forks* over A if $\varphi(x; b) \vdash \bigvee_{i<k} \psi_i(x; c^i)$ and each $\psi_i(x; c^i)$ Kim-divides over A. A type Kim-forks if it implies a formula which does. If $\mathrm{tp}(a/Ab)$ does not Kim-fork over A, we write $a \mathop{\smile\!\!\!\!\!\!\,|}_A^K b$.

Theorem 7.2.3. *[Theorem III.8.1] The following are equivalent for the complete theory T:*

1. *T is NSOP$_1$*

2. *Kim's lemma for Kim-dividing: For every model $M \models T$ and $\varphi(x; b)$, if $\varphi(x; y)$ q-divides for some global M-invariant $q \supseteq tp(b/M)$, then $\varphi(x; y)$ q-divides for* every *global M-invariant $q \supseteq tp(b/M)$.*

3. *Symmetry over models: for every $M \models T$, then $a \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_M^K b$ if and only if $b \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_M^K a$.*

4. *Independence theorem over models: if $M \models T$, $a \equiv_M a'$, $a \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_M^K b$, $a' \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_M^K c$, and $b \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_M^K c$, then there is a'' with $a'' \equiv_{Mb} a$, $a'' \equiv_{Mc} a'$ and $a'' \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_M^K bc$.*

The following definition of an independence relation will be our candidate for Kim-independence in a PAC field with NSOP$_1$ absolute Galois group. It is the conjunction of non-forking indepenence in the ambient separably closed field and Kim-independence in the sense of the Galois group.

Definition 7.2.4. Suppose F is a field and $\mathrm{Th}(\mathcal{S}(\mathcal{G}(F)))$ is NSOP$_1$. If a, b are tuples in some elementary extension of F, we say a and b are *weakly independent* over F if, letting $A = \mathrm{acl}(aF)$ and $B = \mathrm{acl}(bF)$,

1. $A \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_F^{SCF} B$

2. $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_{\mathcal{S}(\mathcal{G}(F))}^K \mathcal{S}(\mathcal{G}(B))$.

Remark 7.2.5. This differs from the definition given by Chatzidakis [Cha17] in the context of $\mathrm{Th}(S(\mathcal{G}(F)))$ having simple theory, as we use Kim-independence and Chatzidakis used non-forking independence in the definition of weak independence. If the base is a model, these definitions agree, and the definition we give here is well-behaved in the broader context of NSOP$_1$ Galois groups.

Theorem 7.2.6. *Suppose F is a PAC field and $Th(S(\mathcal{G}(F)))$ is NSOP$_1$. Then given $L \models Th(F)$, $a \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_L^K b$ if and only if a and b are weakly independent over L.*

Proof. Let $A = \mathrm{acl}(aL)$ and $B = \mathrm{acl}(bL)$.

First, assume $A \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_L^K B$ and we will show A and B are weakly independent over L. Let $(B_i)_{i<\omega}$ be a Morley sequence in a global type finitely satisfiable in L extending $\mathrm{tp}(B/L)$. As $A \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_M^K B$, we may assume $(B_i)_{i<\omega}$ is A-indiscernible. Then by Kim's lemma in the stable theory SCF, we know $A \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_L^{\mathrm{SCF}} B$. Also $(\mathcal{S}(\mathcal{G}(B_i)))_{i<\omega}$ is a Morley sequence in a global type finitely satisfiable in $\mathcal{S}(\mathcal{G}(L))$ which is moreover $\mathcal{S}(\mathcal{G}(A))$-indiscernible. This implies $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_{\mathcal{S}(\mathcal{G}(L))}^K \mathcal{S}(\mathcal{G}(B))$. As $\mathrm{Th}(\mathcal{S}(\mathcal{G}(F))$ is NSOP$_1$, this implies $\mathcal{S}(\mathcal{G}(A)) \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_{\mathcal{S}(\mathcal{G}(F))}^K \mathcal{S}(\mathcal{G}(B))$ by Kim's lemma for Kim-dividing. In other words, A and B are weakly independent over L.

Now assume A and B are weakly independent over L and we will show $A \mathop{\smile\hskip-0.9em\raise0.8ex\hbox{$\mid$}}_L^K B$. Let $(B_i)_{i<\omega}$ be an L-finitely satisfiable Morley sequence over L with $B_0 = B$. Put $C_{n,0} =$

$\mathrm{acl}(B_0, \ldots, B_{2^n-1})$ and $C_{n,1} = \mathrm{acl}(B_{2^n}, \ldots, B_{2^{n+1}-1})$. Note that $C_{n,1} \mathop{\smile\hspace{-0.9em}|}^{u}_{L} C_{n,0}$ and $C_{n,0} \equiv_L C_{n,1}$ for all $n < \omega$. By induction on n, we will choose A_n so that $A_0 = A$ and for all $n \geq 1$ the following conditions are satisfied:

1. $A_{n+1} \equiv_{C_{n,0},C_{n,1}} A_n$.

2. A_n is weakly independent from $\mathrm{acl}(C_{n,0}, C_{n,1}) = C_{n+1,0}$ over L.

3. $A_n C_{n,0} \equiv_L A_n C_{n,1}$.

Suppose we are at stage $n \geq 1$ of the induction. Pick A_n' so that $A_n C_{n+1,0} \equiv_L A_n' C_{n+1,1}$. Then, in particular, $S(\mathcal{G}(A_n)) \equiv_{S(\mathcal{G}(L))} S(\mathcal{G}(A_n'))$, $S(\mathcal{G}(A_n)) \mathop{\smile\hspace{-0.9em}|}^{K}_{S(\mathcal{G}(L))} S(\mathcal{G}(C_{n+1,0}))$, and $S(\mathcal{G}(A_n')) \mathop{\smile\hspace{-0.9em}|}^{K}_{S(\mathcal{G}(L))} S$ Moreover, because $C_{n+1,1} \mathop{\smile\hspace{-0.9em}|}^{u}_{L} C_{n+1,0}$, we have $S(\mathcal{G}(C_{n+1,1})) \mathop{\smile\hspace{-0.9em}|}^{K}_{S(\mathcal{G}(L))} S(\mathcal{G}(C_{n+1,0}))$. As $\mathrm{Th}(S(\mathcal{G}(F)))$ is NSOP_1, we may apply the independence theorem for $\mathop{\smile\hspace{-0.9em}|}^{K}$ to obtain S so that $S \equiv_{S(\mathcal{G}(C_{n+1,0}))} S(\mathcal{G}(A_n))$, $S \equiv_{S(\mathcal{G}(C_{n+1,1}))} S(\mathcal{G}(A_n'))$ and $S \mathop{\smile\hspace{-0.9em}|}^{K}_{S(\mathcal{G}(L))} S(\mathcal{G}(C_{n+1,0}))S(\mathcal{G}(C_{n+1,1}))$. Note that $S(\mathcal{G}(C_{n+2,0})) \supseteq S(\mathcal{G}(C_{n+1,0}))S(\mathcal{G}(C_{n+1,1}))$ so, by extension, we may assume S has been chosen so that $S \mathop{\smile\hspace{-0.9em}|}^{K}_{S(\mathcal{G}(L))} S(\mathcal{G}(C_{n+2,0}))$. By Theorem 7.1.9, there is A_{n+1} with $A_{n+1} \equiv_{C_{n+1,0}} A_n$, $A_{n+1} \equiv_{C_{n+1,1}} A_n'$, $A_{n+1} \mathop{\smile\hspace{-0.9em}|}^{\mathrm{SCF}}_{L} C_{n+1,0}C_{n+1,1}$ and $S(\mathcal{G}(A_{n+1})) = S$. Note that this implies $A_{n+1} \mathop{\smile\hspace{-0.9em}|}^{\mathrm{SCF}}_{L} C_{n+2,0}$ because $C_{n+2,0} = \mathrm{acl}(C_{n+1,0}C_{n+1,1})$ and $S(\mathcal{G}(A_{n+1})) \mathop{\smile\hspace{-0.9em}|}^{K}_{S(\mathcal{G}(L))} S(\mathcal{G}(C_{n+2,0}))$ by the choice of S so A_{n+1} and $C_{n+2,0}$ are weakly independent over L. This completes the induction.

Let $p(X; B) = \mathrm{tp}(A/B)$. Observe that, by construction, $A_n \models \bigcup_{i<2^n} p(X; B_i)$ so $\bigcup_{i<\omega} p(X; B_i)$ is consistent by compactness. As the sequence $\langle B_i : i < \omega \rangle$ is an arbitrary L-finitely satisfiable Morley sequence over L, this shows $A \mathop{\smile\hspace{-0.9em}|}^{K}_{L} B$. $\square$

Corollary 7.2.7. *If F is a PAC field such that $Th(S(\mathcal{G}(F)))$ is $NSOP_1$, then $Th(F)$ is $NSOP_1$.*

Proof. A theory is NSOP_1 if and only if $\mathop{\smile\hspace{-0.9em}|}^{K}$ is symmetric by Theorem 3.8.1. If $\mathrm{Th}(S(\mathcal{G}(F)))$ is NSOP_1, then $\mathop{\smile\hspace{-0.9em}|}^{K}$ in $\mathrm{Th}(S(\mathcal{G}(F)))$ is symmetric and $\mathop{\smile\hspace{-0.9em}|}^{\mathrm{SCF}}$ is symmetric by stability of the theory of separably closed fields. Therefore $\mathop{\smile\hspace{-0.9em}|}^{K}$ is symmetric in $\mathrm{Th}(F)$ by Theorem 7.2.6. $\square$

SOP_2

In the SOP_2 case, we do not have a theory of independence as in the SOP_1 case, but we can find amalgamation problems from strongly indiscernible trees. The basic principle we will use is that if we have a complete type $p(x; b)$ and a strongly indiscernible tree $(b_\eta)_{\eta \in 2^{<\omega}}$, then if $\bigcup_{n<\omega} p(x; b_{0^n})$ is consistent, then $p(x; b_\eta) \cup p(x; b_\nu)$ is consistent for any two $\eta \perp \nu \in 2^{<\omega}$. Otherwise, by compactness and indiscernibility there would be a formula in p witnessing SOP_2. This weak form of amalgamation in an NSOP_2 absolute Galois group suffices to prove that the field is also NSOP_2.

Proposition 7.2.8. *Let F be a PAC field. Then $Th(F)$ is NSOP$_2$ if $Th(\mathcal{S}(\mathcal{G}(F)))$ is NSOP$_2$.*

Proof. We may suppose F is a monster model of its theory and fix a Skolemization of F. Then, formally, we work in F^s with a predicate for F and these Skolem functions—in the proof, indiscernbility will always mean with respect to this structure, though acl means model-theoretic algebraic closure in F, as a pure field. Towards contradiction, suppose F has SOP$_2$ witnessed by some formula $\varphi(x; y)$ in the language of rings. By compactness, we may suppose $\varphi(x; y)$ witnesses SOP$_2$ via the parameters $(b_\eta)_{\eta \in 2^{<\omega+\omega}}$, which form a strongly indiscernible tree. In F, choose a so that $a \models \{\varphi(x; b_{0^\alpha}) : \alpha < \omega + \omega\}$. By Ramsey, compactness, and automorphism, we may assume $(b_{0^\alpha})_{\alpha<\omega+\omega}$ is a-indiscernible.

Let E be the Skolem hull of $b_{<\omega}$ in F. Define $A = \mathrm{acl}(aE)$ and set $B_\eta = \mathrm{acl}(Eb_{0^\omega \frown \eta})$ for $\eta \in 2^{<\omega}$. Then $(B_\eta)_{\eta\in 2^{<\omega}}$ is strongly indiscernible over E, $(B_{0^\alpha})_{\alpha<\omega}$ is A-indiscernible, and $(B_{0^\alpha})_{\alpha<\omega}$ is an E-finitely satisfiable Morley sequence, enumerated in reverse. Then by Kim's lemma in the stable theory SCF, we have $A \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\mathrm{SCF}}_E B_0$. By strong indiscernibility, $(B_{0^\alpha})_{\alpha<\omega}$ is also $B_{\langle 1\rangle}$-indiscernible, so also $B_0 \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\mathrm{SCF}}_E B_{\langle 1\rangle}$. Choose A' so that $A'B_{\langle 1\rangle} \equiv_E AB_0$.

Let $q(X; S(\mathcal{G}(B_\emptyset))) = \mathrm{tp}(S(\mathcal{G}(A))/S(\mathcal{G}(B_\emptyset)))$. We know $(S(\mathcal{G}(B_\eta)))_{\eta\in 2^{<\omega}}$ is a strongly indiscernble tree over $S(\mathcal{G}(E))$ and $S(\mathcal{G}(A)) \models \bigcup_{n<\omega} q(X; S(\mathcal{G}(B_{0^n})))$ as $(S(\mathcal{G}(B_{0^n})))_{n<\omega}$ is $S(\mathcal{G}(A))$-indiscernible. As $\mathrm{Th}(S(\mathcal{G}(F)))$ is NSOP$_2$, $q(X; B_0) \cup q(X; B_{\langle 1\rangle})$ is consistent. Let S_0 be a realization. Then by Theorem 7.1.9, there is A_* with $A_* \equiv_{B_0} A$, $A_* \equiv_{B_{\langle 1\rangle}} A'$, and $S(\mathcal{G}(A_*)) = S_0$. By the choice of A and definition of B_0 and $B_{\langle 1\rangle}$, this implies $\{\varphi(x; b_{0^\omega \frown 0}), \varphi(x; b_{0^\omega \frown \langle 1\rangle})\}$ is consistent. This contradicts the definition of SOP$_2$, completing the proof. $\square$

Bibliography

[Adl05] Hans Adler. Explanation of independence. *arXiv preprint math/0511616*, 2005.

[Adl14] Hans Adler. Kim's lemma for ntp2 theories: a simpler proof of a result by chernikov and kaplan. 2014.

[Bau02] Andreas Baudisch. Generic variations of models of t. *The Journal of Symbolic Logic*, 67(3):1025–1038, 2002.

[BKM78] Jon Barwise, Matt Kaufmann, and Michael Makkai. Stationary logic. *Annals of Mathematical Logic*, 13(2):171–224, 1978.

[C$^+$17] Gabriel Conant et al. Forking and dividing in henson graphs. *Notre Dame Journal of Formal Logic*, 58(4):555–566, 2017.

[Cas11] Enrique Casanovas. *Simple theories and hyperimaginaries*, volume 39. Cambridge University Press, 2011.

[CH03] Gregory L Cherlin and Ehud Hrushovski. *Finite structures with few types*. Number 152. Princeton University Press, 2003.

[Cha98] Zoé Chatzidakis. Model theory of profinite groups having the iwasawa property. *Illinois Journal of Mathematics*, 42(1):70–96, 1998.

[Cha99] Zoé Chatzidakis. Simplicity and independence for pseudo-algebraically closed fields. In *Models and computability (Leeds, 1997)*, volume 259 of *London Math. Soc. Lecture Note Ser.*, pages 41–61. Cambridge Univ. Press, Cambridge, 1999.

[Cha02] Zoé Chatzidakis. Properties of forking in ω-free pseudo-algebraically closed fields. *The Journal of Symbolic Logic*, 67(03):957–996, 2002.

[Cha17] Zoé Chatzidakis. Amalgamation of types in pseudo-algebraically closed fields and applications. *arXiv preprint arXiv:1704.08132*, 2017.

[Che14] Artem Chernikov. Theories without the tree property of the second kind. *Ann. Pure Appl. Logic*, 165(2):695–723, 2014.

[CHL85] Gregory Cherlin, Leo Harrington, and Alistair H Lachlan. âĎţ0-categorical, âĎţ0-stable structures. *Annals of Pure and Applied Logic*, 28(2):103–135, 1985.

[CK90] Chen Chung Chang and H Jerome Keisler. *Model theory*, volume 73. Elsevier, 1990.

[CK12] Artem Chernikov and Itay Kaplan. Forking and dividing in NTP_2 theories. *The Journal of Symbolic Logic*, 77(01):1–20, 2012.

[Con] Gabriel Conant. "Map of the universe". `http://www.forkinganddividing.com/`.

[Con14] Gabriel Conant. Forking and dividing in henson graphs. *arXiv preprint arXiv:1401.1570*, 2014.

[CP98] Zoé Chatzidakis and Anand Pillay. Generic structures and simple theories. *Annals of Pure and Applied Logic*, 95(1):71–92, 1998.

[CR16] Artem Chernikov and Nicholas Ramsey. On model-theoretic tree properties. *Journal of Mathematical Logic*, page 1650009, 2016.

[CvdDM80] Gregory Cherlin, Lou van den Dries, and Angus Macintyre. The elementary theory of regularly closed fields. *preprint*, 1980.

[DS04] Mirna Džamonja and Saharon Shelah. On $\triangleleft^*$-maximality. *Annals of Pure and Applied Logic*, 125(1):119–158, 2004.

[FJ08] Michael D Fried and Moshe Jarden. Field arithmetic. a series of modern surveys in mathematics, 11, 2008.

[Gal80] Fred Galvin. Chain conditions and products. *Fund. Math.*, 108(1):33–48, 1980.

[GIL02] Rami Grossberg, José Iovino, and Olivier Lessmann. A primer of simple theories. *Archive for Mathematical Logic*, 41(6):541–580, 2002.

[Gra99] Nicholas Granger. *Stability, simplicity, and the model theory of bilinear forms*. PhD thesis, University of Manchester, 1999.

[Hod93] Wilfrid Hodges. *Model theory*, volume 42. Cambridge University Press Cambridge, 1993.

[HP94] Ehud Hrushovski and Anand Pillay. Groups definable in local fields and pseudo-finite fields. *Israel Journal of Mathematics*, 85(1):203–262, 1994.

[Hru91] Ehud Hrushovski. Pseudo-finite fields and related structures. *Model theory and applications*, 11:151–212, 1991.

[Hru97] Ehud Hrushovski. Simplicity and the lascar group. *preprint*, 231:232–235, 1997.

[Hru12] Ehud Hrushovski. Stable group theory and approximate subgroups. *Journal of the American Mathematical Society*, 25(1):189–243, 2012.

[HS14] Gwyneth Harrison-Shermoen. *Independence Relations in Theories with the Tree Property*. PhD thesis, UC Berkeley, 2014.

[Jec13] Thomas Jech. *Set theory*. Springer Science & Business Media, 2013.

[Jeř17] Emil Jeřábek. Recursive functions and existentially closed structures. *arXiv preprint arXiv:1710.09864*, 2017.

[Kan03] Akihiro Kanamori. *The higher infinite*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, second edition, 2003. Large cardinals in set theory from their beginnings.

[Kim98] Byunghan Kim. Forking in simple unstable theories. *Journal of the London Mathematical Society*, 57(02):257–267, 1998.

[Kim01] Byunghan Kim. Simplicity, and stability in there. *The Journal of Symbolic Logic*, 66(02):822–836, 2001.

[Kim09] Byunghan Kim. NTP_1 theories. Slides, Stability Theoretic Methods in Unstable Theories, BIRS, February 2009.

[KK11] Byunghan Kim and Hyeung-Joon Kim. Notions around tree property 1. *Ann. Pure Appl. Logic*, 162(9):698–709, 2011.

[KKS14] Byunghan Kim, Hyeung-Joon Kim, and Lynn Scow. Tree indiscernibilities, revisited. *Arch. Math. Logic*, 53(1-2):211–232, 2014.

[KLM] WM KANTOR, MW LIEBECK, and HD MACPHERSON. K0-categorical structures smoothly approximated by finite substructures. *Proceedings of the London Mathematical Society*, pages 439–463.

[KLM89] William M Kantor, Martin W Liebeck, and H Dugald Macpherson. $\aleph_0$-categorical structures smoothly approximated by finite substructures. *Proceedings of the London Mathematical Society*, 3(3):439–463, 1989.

[KLS16] Itay Kaplan, Noa Lavi, and Saharon Shelah. The generic pair conjecture for dependent finite diagrams. *Israel J. Math.*, 212(1):259–287, 2016.

[KP97] Byunghan Kim and Anand Pillay. Simple theories. *Annals of Pure and Applied Logic*, 88(2-3):149–164, 1997.

[KPT05] Alexander S Kechris, Vladimir G Pestov, and Stevo Todorcevic. Fraïssé limits, ramsey theory, and topological dynamics of automorphism groups. *Geometric and Functional Analysis*, 15(1):106–189, 2005.

[KR17] Itay Kaplan and Nicholas Ramsey. On kim-independence. *J. Eur. Math. Soc. (JEMS)*, 2017. accepted, `arXiv:1702.03894`.

[KS02] Juliette Kennedy and Saharon Shelah. On regular reduced products. *J. Symbolic Logic*, 67(3):1169–1177, 2002.

[Kue72] David W Kueker. Löwenheim-skolem and interpolation theorems in infinitary languages. *Bulletin of the American Mathematical Society*, 78(2):211–215, 1972.

[Lac75] Alistair H Lachlan. A remark on the strict order property. *Mathematical Logic Quarterly*, 21(1):69–70, 1975.

[LMT09] Martin W Liebeck, Dugald Macpherson, and Katrin Tent. Primitive permutation groups of bounded orbital diameter. *Proceedings of the London Mathematical Society*, 100(1):216–248, 2009.

[Mal12] ME Malliaris. Hypergraph sequences as a tool for saturation of ultrapowers. *The Journal of Symbolic Logic*, 77(01):195–223, 2012.

[Med] Alice Medvedev. Qacfa. *Preprint, available on http://arxiv.org/abs/1508.06007*.

[MS86] Alan H Mekler and Saharon Shelah. Stationary logic and its friends. ii. *Notre Dame Journal of Formal Logic*, 27(1):39–50, 1986.

[MS15a] M Malliaris and S Shelah. Model-theoretic applications of cofinality spectrum problems. *arXiv preprint arXiv:1503.08338*, 2015.

[MS15b] Maryanthe Malliaris and Saharon Shelah. Constructing regular ultrafilters from a model-theoretic point of view. *Transactions of the American Mathematical Society*, 367(11):8139–8173, 2015.

[MS16] Maryanthe Malliaris and Saharon Shelah. Cofinality spectrum theorems in model theory, set theory, and general topology. *Journal of the American Mathematical Society*, 29(1):237–297, 2016.

[Nüb04] Herwig Nübling. Adding skolem functions to simple theories. *Archive for Mathematical Logic*, 43(3):359–370, 2004.

[Sco12] L. Scow. Indiscernibles, EM-types, and Ramsey Classes of Trees. *ArXiv e-prints*, August 2012.

[She75] Saharon Shelah. Generalized quantifiers and compact logic. *Transactions of the American Mathematical Society*, 204:342–364, 1975.

[She80] Saharon Shelah. Simple unstable theories. *Annals of Mathematical Logic*, 19(3):177–203, 1980.

[She90] Saharon Shelah. *Classification theory: and the number of non-isomorphic models*. Elsevier, 1990.

[She94] Saharon Shelah. *Cardinal Arithmetic*, volume 29 of *Oxford Logic Guides*. Oxford University Press, 1994.

[She97] Saharon Shelah. Colouring and non-productivity of $\aleph_2$-cc. *Annals of Pure and Applied Logic*, 84:153–174, 1997. arxiv:math.LO/9609218.

[Sim15] Pierre Simon. *A guide to NIP theories*. Cambridge University Press, 2015.

[SU08] Saharon Shelah and Alexander Usvyatsov. More on sop 1 and sop 2. *Annals of Pure and Applied Logic*, 155(1):16–31, 2008.

[TT12] Kota Takeuchi and Akito Tsuboi. On the existence of indiscernible trees. *Ann. Pure Appl. Logic*, 163(12):1891–1902, 2012.

[Win75] Peter Winkler. Model-completeness and skolem expansions. *Model Theory and Algebra*, pages 408–463, 1975.

[YC14] Itaï Ben Yaacov and Artem Chernikov. An independence theorem for ntp 2 theories. *The Journal of Symbolic Logic*, 79(01):135–153, 2014.