A Unified Approach to Structural Limits and Limits of Graphs with Bounded Tree-depth

Jaroslav Nešetřil

Patrice Ossona de Mendez

Author address:

Jaroslav Nešetřil, Computer Science Institute of Charles University (IUUK and ITI), Malostranské nám.25, 11800 Praha 1, Czech Republic
E-mail address: nesetril@iuuk.mff.cuni.cz

Patrice Ossona de Mendez, Centre d'Analyse et de Mathématiques Sociales (CNRS, UMR 8557), 190-198 avenue de France, 75013 Paris, France
E-mail address: pom@ehess.fr

Contents

Abstract

In this paper we introduce a general framework for the study of limits of relational structures and graphs in particular, which is based on a combination of model theory and (functional) analysis. We show how the various approaches to graph limits fit to this framework and that they naturally appear as "tractable cases" of a general theory. As an outcome of this, we provide extensions of known results. We believe that this put these into next context and perspective. For example, we prove that the sparse–dense dichotomy exactly corresponds to random free graphons. The second part of the paper is devoted to the study of sparse structures. First, we consider limits of structures with bounded diameter connected components and we prove that in this case the convergence can be "almost" studied component-wise. We also propose the structure of limits objects for convergent sequences of sparse structures. Eventually, we consider the specific case of limits of colored rooted trees with bounded height and of graphs with bounded tree-depth, motivated by their role as "elementary bricks" these graphs play in decompositions of sparse graphs, and give an explicit construction of a limit object in this case. This limit object is a graph built on a standard probability space with the property that every first-order definable set of tuples is measurable. This is an example of the general concept of *modeling* we introduce here. Our example is also the first "intermediate class" with explicitly defined limit structures where the inverse problem has been solved.

Received by the editor October 20th, 2013.

2010 *Mathematics Subject Classification.* Primary 03C13 (Finite structures), 03C98 (Applications of model theory), 05C99 (Graph theory), 06E15 (Stone spaces and related structures), Secondary 28C05 (Integration theory via linear functionals).

Key words and phrases. Graph and Relational structure and Graph limits and Structural limits and Radon measures and Stone space and Model theory and First-order logic and Measurable graph.

Supported by grant ERCCZ LL-1201 and CE-ITI P202/12/G061, and by the European Associated Laboratory "Structures in Combinatorics" (LEA STRUCO).

Supported by grant ERCCZ LL-1201 and by the European Associated Laboratory "Structures in Combinatorics" (LEA STRUCO).

CHAPTER 1

Introduction

To facilitate the study of the asymptotic properties of finite graphs (and more generally of finite structures) in a sequence $G_1, G_2, \dots, G_n, \dots$, it is natural to introduce notions of *structural convergence*. By structural convergence, we mean that we are interested in the characteristics of a typical vertex (or group of vertices) in the graph G_n, as n grows to infinity. This convergence can be concisely expressed by various means. We note two main directions:

- the convergence of the sampling distributions;
- the convergence with respect to a metric in the space of structures (such as the cut metric).

Also, sampling from a limit structure may also be used to define a sequence convergent to the limit structure.

All these directions lead to a rich theory which originated in a probabilistic context by Aldous [3] and Hoover [42] (see also the monograph of Kallenberg [45] and the survey of Austin [6]) and, independently, in the study of random graph processes, and in analysis of properties of random (and quasirandom) graphs (in turn motivated among others by statistical physics [13, 14, 54]). This development is nicely documented in the recent monograph of Lovász [53].

The asymptotic properties of large graphs are studied also in the context of decision problems as exemplified e.g. by structural graphs theory, [22, 71]. However it seems that the existential approach typical for decision problems, structural graph theory and model theory on the one side and the counting approach typical for statistics and probabilistic approach on the other side have little in common and lead to different directions: on the one side to study, say, definability of various classes and the properties of the homomorphism order and on the other side, say, properties of partition functions. It has been repeatedly stated that these two extremes are somehow incompatible and lead to different area of study (see e.g. [11, 40]). In this paper we take a radically different approach which unifies these both extremes.

We propose here a model which is a mixture of the analytic, model theoretic and algebraic approach. It is also a mixture of existential and probabilistic approach. Precisely, our approach is based on the *Stone pairing* $\langle \phi, G\rangle$ of a first-order formula ϕ (with set of free variables $\mathrm{Fv}(\phi)$) and a graph G, which is defined by following expression

$$\langle \phi, G\rangle = \frac{|\{(v_1,\dots,v_{|\mathrm{Fv}(\phi)|}) \in G^{|\mathrm{Fv}(\phi)|} : G \models \phi(v_1,\dots,v_{|\mathrm{Fv}(\phi)|})\}|}{|G|^{|\mathrm{Fv}(\phi)|}}.$$

Stone pairing induces a notion of convergence: a sequence of graphs $(G_n)_{n\in\mathbb{N}}$ is FO-*convergent* if, for every first order formula ϕ (in the language of graphs), the values $\langle \phi, G_n\rangle$ converge as $n \to \infty$. In other words, $(G_n)_{n\in\mathbb{N}}$ is FO-convergent if the probability that a formula ϕ is satisfied by the graph G_n with a random assignment of vertices of G_n to the free variables of ϕ converges as n grows to infinity. We also consider analogously defined X-convergence, where X is a fragment of FO.

Our main result is that this model of FO-convergence is a suitable model for the analysis of limits of sparse graphs (and particularly of graphs with bounded tree depth). This fits to a broad context of recent research.

For graphs, and more generally for finite structures, there is a class dichotomy: *nowhere dense* and *somewhere dense* [69, 65]. Each class of graphs falls in one of these two categories. Somewhere dense class $\mathcal{C}$ may be characterised by saying that there exists a (primitive positive) FO interpretation of all graphs into them. Such class $\mathcal{C}$ is inherently a class of dense graphs. In the theory of nowhere dense structures [71] there are two extreme conditions related to sparsity: bounded degree and bounded diameter. Limits of bounded degree graphs have been studied thoroughly [8], and this setting has been partialy extended to sparse graphs with far away large degree vertices [56]. The class of graphs with bounded diameter is considered in Section 3.3 (and leads to a difficult analysis of componentwise convergence). This analysis provides a first-step for the study of limites of graphs with bounded *tree-depth*. Classes of graphs with bounded tree-depth can be defined by logical terms as well as combinatorially in various ways; the most concise definition is perhaps that a class of graphs has bounded tree depth if and only if the maximal length of a path in every G in the class is bounded by a constant. Graphs with bounded tree-depth play also the role of building blocks of graphs in a nowhere dense class (by means of *low tree-depth decompositions* [59, 60, 71]). So the solution of limits for graphs with bounded tree depth presents a step (and perhaps provides a road map) in solving the limit problem for sparse graphs.

We propose here a new type of measurable structure, called *modeling*, which extends the notion of graphing, and which we believe is a good candidate for limit objects of sequence of graphs in a nowhere dense class. The convergence of graphs with bounded tree depth is analysed in detail and this leads to a construction of a modeling limits for those sequences of graphs where all members of the sequence have uniformly bounded tree depth (see Theorem 4.36). Moreover, we characterize modelings which are limits of graphs with bounded tree-depth.

There is more to this than meets the eye: We prove that if $\mathcal{C}$ is a monotone class of graphs such that every FO-convergent sequence has a modeling limit then the class $\mathcal{C}$ is nowhere dense (see Theorem 3.32). This shows the natural limitations to modeling FO-limits. To create a proper model for bounded height trees we have to introduce the model in a greater generality and it appeared that our approach relates and in most cases generalizes, by properly chosing fragment X of FO, all existing models of graph limits. For instance, for the fragment X of all existential first-order formulas, X-convergence means that the probability that a structure has a particular extension property converges. Our approach is encouraged by the deep connections to the *four notions* of convergence which have been proposed to study graph limits in different contexts.

The ultimate goal of the study of structural limits is to provide (as effectively as possible) limit objects themselves: we would like to find an object which will induce the limit distribution and encode the convergence. This was done in a few isolated cases only: For dense graphs Lovász and Szegedy isolated the notion of *graphon*: In this representation the limit [54, 13] is a symmetric Lebesgue measurable function $W : [0,1]^2 \to [0,1]$ called *graphon*.

A representation of the limit (for our second example of bounded degree graphs) is a *measurable graphing* (notion introduced by Adams [1] in the context of Ergodic theory), that is a standard Borel space with a measure μ and d measure preserving Borel involutions. The existence of such a representation has been made explicit by Elek [28], and relies on the works of Benjamini [8] and Gaboriau [34]. Both of these models of convergence are particular cases of our general approach

One of the main issue of our general approach is to determine a representation of FO-limits as measurable graphs. A natural limit object is a standard probability space (V, Σ, μ) together with a graph with vertex set V and edge set E, with the property that every first-order definable subset of a power of V is measurable. This leads to the notion of relational sample space and to the notion of *modelling*. This notion seems to be particularly suitable for sparse graphs (and in the full generality only for sparse graphs, see Theorem 3.32.).

In this paper, we shed a new light on all these constructions by an approach inspired by functional analysis. The preliminary material and our framework are introduced in Sections 1.1 and 2.1. The general approach presented in the first sections of this paper leads to several new results. Let us mention a sample of such results.

Central to the theory of graph limits stand random graphs (in the Erdős-Rényi model [30]): a sequence of random graphs with increasing order and edge probability $0 < p < 1$ is almost surely convergent to the constant graphon p [54]. On the other hand, it follows from the work of Erdős and Rényi [31] that such a sequence is almost surely elementarily convergent to an ultra-homogeneous graph, called the *Rado graph*. We prove that these two facts, together with the quantifier elimination property of ultra-homogeneous graphs, imply that a sequence of random graphs with increasing order and edge probability $0 < p < 1$ is almost surely FO-convergent, see Section 2.3.4. (However, we know that this limit cannot be neither random free graphon nor modelling, see Theorem 3.32)

We shall prove that a sequence of bounded degree graphs $(G_n)_{n\in\mathbb{N}}$ with $|G_n| \to \infty$ is FO-convergent if and only if it is both convergent in the sense of Benjamini-Schramm and in the sense of elementary convergence. The limit can still be represented by a graphing, see Sections 2.2.2 and 3.2.6.

Why Stone pairing? We prove that the limit of an FO-convergent sequence of graphs is a probability measure on the Stone space of the Boolean algebra of first-order formulas, which is invariant under the action of S_ω on this space, see Section 2.1. Fine interplay of these notions is depicted on Table 1.

Graph limits (in the sense of Lovász et al.) — and more generally hypergraph limits — have been studied by Elek and Szegedy [29] through the introduction of a measure on the ultraproduct of the graphs in the sequence (via Loeb measure construction, see [50]). The fundamental theorem of ultraproducts proved by Łoś [51] implies that the ultralimit of a sequence of graphs is (as a measurable graph) an FO-limit. Thus in this non-standard setting we get FO-limits (almost) for free see [70]. (However this does not relate to the difficult question of separability.)

We believe that the approach taken in this paper is natural and that it enriches the existing notions of limits. In a sense we proceed dually to [11]: We do not view $\langle \phi, G\rangle$ as a "ϕ test" for G but rather as operator induced by G on the Boolean

Boolean algebra $\mathcal{B}(X)$	Stone Space $S(\mathcal{B}(X))$
Formula ϕ	Continuous function f_ϕ
Vertex v	"Type of vertex" T
Graph G	statistics of types =probability measure μ_G
$\langle \phi, G\rangle$	$\int f_\phi(T)\ \mathrm{d}\mu_G(T)$
X-convergent (G_n)	weakly convergent μ_{G_n}
$\Gamma = \mathrm{Aut}(\mathcal{B}(X))$	Γ-invariant measure

TABLE 1. Some correspondances

algebra of all FO-formulas (or on the subalgebra induced by a fragment $X \subset \mathrm{FO}$). It also presents, for example via decomposition techniques (low-tree depth decomposition, see [71]) a promising approach to more general intermediate classes (see the final comments).

1.1. Main Definitions and Results

If we consider relational structures with signature λ, the symbols of the relations and constants in λ define the non-logical symbols of the vocabulary of the first-order language $\mathrm{FO}(\lambda)$ associated to λ-structures. Notice that if λ is at most countable then $\mathrm{FO}(\lambda)$ is countable. The symbols of variables will be assumed to be taken from a countable set $\{x_1, \dots, x_n, \dots\}$ indexed by $\mathbb{N}$. Let $u_1, \dots, u_k$ be terms. The set of used free variables of a formula ϕ will be denoted by $\mathrm{Fv}(\phi)$ (by saying that a variable x_i is "used" in ϕ we mean that ϕ is not logically equivalent to a formula in which x_i does not appear). The formula $\phi_{x_{i_1},\dots,x_{i_k}}(u_1, \dots, u_k)$ denote the formula obtained by substituting simultaneously the term u_j to the free occurences of x_{i_j} for $j = 1, \dots, k$. In the sake of simplicity, we will denote by $\phi(u_1, \dots, u_k)$ the substitution $\phi_{x_1,\dots,x_k}(u_1, \dots, u_k)$.

A *relational structure* $\mathbf{A}$ with signature λ is defined by its *domain* (or *universe*) A and relations with names and arities as defined in λ. In the following we will denote relational structures by bold face letters $\mathbf{A}, \mathbf{B}, \dots$ and their domains by the corresponding light face letters $A, B, \dots$

The key to our approach are the following two definitions.

DEFINITION 1.1 (Stone pairing). Let λ be a signature, let $\phi \in \mathrm{FO}(\lambda)$ be a first-order formula with free variables $x_1, \dots, x_p$ and let $\mathbf{A}$ be a finite λ-structure.

Put

$$\Omega_\phi(\mathbf{A}) = \{(v_1, \dots, v_p) \in A^p :\ \mathbf{A} \models \phi(v_1, \dots, v_p)\}.$$

We define the *Stone pairing* of ϕ and $\mathbf{A}$ by

$$\langle \phi, \mathbf{A}\rangle = \frac{|\Omega_\phi(\mathbf{A})|}{|A|^p}. \tag{1.1}$$

In other words, $\langle \phi, \mathbf{A}\rangle$ is the probability that ϕ is satisfied in $\mathbf{A}$ when we interpret the p free variables of ϕ by p vertices of G chosen randomly, uniformly and independently. Also, $\Omega_\phi(\mathbf{A})$ is interpreted as the *solution set* of ϕ in $\mathbf{A}$.

Note that in the case of a sentence ϕ (that is a formula with no free variables, thus $p = 0$), the definition of the Stone pairing reduces to

$$\langle \phi, \mathbf{A}\rangle = \begin{cases} 1, & \text{if } \mathbf{A} \models \phi; \\ 0, & \text{otherwise.} \end{cases}$$

Definition 1.2 (FO-convergence). A sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of finite λ-structures is FO-*convergent* if, for every formula $\phi \in \mathrm{FO}(\lambda)$ the sequence $(\langle \phi, \mathbf{A}_n\rangle)_{n\in\mathbb{N}}$ is (Cauchy) convergent.

In other words, a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is FO-convergent if the sequence of mappings $\langle\,\cdot\,, \mathbf{A}_n\rangle : \mathrm{FO}(\lambda) \to [0, 1]$ is pointwise-convergent.

The interpretation of the Stone pairing as a probability suggests to extend this view to more general λ-structures which will be our candidates for limit objects.

Definition 1.3 (Relational sample space). A *relational sample space* is a relational structure $\mathbf{A}$ (with signature λ) with extra structure: The domain A of $\mathbf{A}$ of a sample model is a standard Borel space (with Borel σ-algebra $\Sigma_{\mathbf{A}}$) with the property that every subset of A^p that is first-order definable in $\mathrm{FO}(\lambda)$ is measurable (in A^p with respect to the product σ-algebra). For brevity we shall use the same letter $\mathbf{A}$ for structure and relational sample space.

In other words, if $\mathbf{A}$ is a relational sample space then for every integer p and every $\phi \in \mathrm{FO}(\lambda)$ with p free variables it holds $\Omega_\phi(\mathbf{A}) \in \Sigma_{\mathbf{A}}^p$.

Definition 1.4 (Modeling). A *modeling* $\mathbf{A}$ is a relational sample space $\mathbf{A}$ equipped with a probability measure (denoted $\nu_{\mathbf{A}}$). By the abuse of symbols the modelling will be denoted by $\mathbf{A}$ (with σ-algebra $\Sigma_{\mathbf{A}}$ and corresponding measure $\nu_{\mathbf{A}}$). A *modeling* with signature λ is a λ-*modeling*.

Remark 1.5. We take time for some comments on the above definitions:

- According to Kuratowski's isomorphism theorem, the domains of relational sample spaces are Borel-isomorphic to either $\mathbb{R}$, $\mathbb{Z}$, or a finite space.
- *Borel graphs* (in the sense of Kechris et al. [46]) are generally not modelings (in our sense) as Borel graphs are only required to have a measurable adjacency relation.
- By equipping its domain with the discrete σ-algebra, every finite λ-structure defines a relational sample space. Considering the uniform probability measure on this space then canonically defines a uniform modeling.

- It follows immediately from Definition 1.3 that any *k-rooting* of a relational sample space is a relational sample space.

We can extend the definition of Stone pairing from finite structures to modelings as follows.

Definition 1.6 (Stone pairing for modeling). Let λ be a signature, let $\phi \in \mathrm{FO}(\lambda)$ be a first-order formula with free variables $x_1, \dots, x_p$ and let $\mathbf{A}$ be a λ-modeling.

We can define the *Stone pairing* of ϕ and $\mathbf{A}$ by

$$\langle \phi, \mathbf{A} \rangle = \int_{x \in A^p} 1_{\Omega_\phi(\mathbf{A})}(x)\, \mathrm{d}\nu_{\mathbf{A}}^p(x). \tag{1.2}$$

Note that the definition of a modeling is simply tailored to make the expression (1.2) meaningful. Based on this definition, modelings can sometimes be used as a representation of the limit of an FO-convergent sequence of finite λ-structures.

Definition 1.7. A modeling $\mathbf{L}$ is a *modeling* FO*-limit* of an FO-convergent sequence $(\mathbf{A}_n)_{n \in \mathbb{N}}$ of finite λ-structures if $\langle \phi\,, \mathbf{A}_n \rangle$ converges pointwise to $\langle\, \phi\,, \mathbf{L} \rangle$ for every first order formula ϕ.

As we shall see in Lemma 3.8, a modeling FO-limit of an FO-convergent sequence $(\mathbf{A}_n)_{n \in \mathbb{N}}$ of finite λ-structures is necessarily weakly uniform. It follows that if a modeling $\mathbf{L}$ is a modeling FO-limit then L is either finite or uncountable.

We shall see that not every FO-convergent sequence of finite relational structures admits a modeling FO-limit. In particular we prove:

Theorem 1.8. *Let $\mathcal{C}$ be a monotone class of finite graphs, such that every* FO*-convergent sequence of graphs in $\mathcal{C}$ has a modeling* FO*-limit. Then the class $\mathcal{C}$ is nowhere dense.*

Recall that a class of graphs is *monotone* if it is closed by the operation of taking a subgraph, and that a monotone class of graphs $\mathcal{C}$ is *nowhere dense* if, for every integer p, there exists an integer $N(p)$ such that the p-th subdivision of the complete graph $K_{N(p)}$ on $N(p)$ vertices does not belong to $\mathcal{C}$ (see [65, 69, 71]).

However, we conjecture that the theorem above expresses exactly when modeling FO-limits exist:

Conjecture 1.1. If $(G_n)_{n \in \mathbb{N}}$ is an FO-convergent sequence of graphs and if $\{G_n : n \in \mathbb{N}\}$ is a nowhere dense class, then the sequence $(G_n)_{n \in \mathbb{N}}$ has a modeling FO-limit.

As a first step, we prove that modeling FO-limits exist in two particular cases, which form in a certain sense the building blocks of nowhere dense classes.

Theorem 1.9. *Let C be a integer.*

(1) Every FO-convergent sequence of graphs with maximum degree at most C has a modeling FO-limit;

(2) *Every FO-convergent sequence of rooted trees with height at most C has a modeling FO-limit.*

The first item will be derived from the graphing representation of limits of Benjamini-Schramm convergent sequences of graphs with bounded maximum degree with no major difficulties. Recall that a *graphing* [1] is a Borel graph G such that the following *Intrinsic Mass Transport Principle* (IMTP) holds:

$$\forall A, B \qquad \int_A \deg_B(x)\,\mathrm{d}x = \int_B \deg_A(y)\mathrm{d}y,$$

where the quantification is on all measurable subsets of vertices, and where $\deg_B(x)$ (resp. $\deg_A(y)$) denote the degree in B (resp. in A) of the vertex x (resp. of the vertex y). In other words, the Mass Transport Principle states that if we count the edges between sets A and B be summing up the degrees in B of vertices in A or by summing up the degrees in A of vertices in B, we should get the same result.

Theorem 1.10 (Elek [28]). *The Benjamini-Schramm limit of a bounded degree graph sequence can be represented by a graphing.*

A full characterization of the limit objects in this case is not known, and is related to the following conjecture.

Conjecture 1.2 (Aldous, Lyons [5]). Every graphing is the Benjamini-Schramm limit of a bounded degree graph sequence.

Equivalently, every unimodular distribution on rooted countable graphs with bouded degree is the Benjamini-Schramm limit of a bounded degree graph sequence.

We conjecture that a similar condition could characterize modeling FO-limits of sequences of graphs with bounded degree. In this more general setting, we have to add a new condition, namely to have the *finite model property.* Recall that an infinite structure $\mathbf{L}$ has the finite model property if every sentence satisfied by $\mathbf{L}$ has a finite model.

Conjecture 1.3. A modeling is the Benjamini-Schramm limit of a bounded degree graph sequence if and only if it is a graph with bounded degree, is weakly uniform, it satisfies both the Intrinsic Mass Transport Principle, and it has the finite model property.

When handling infinite degrees, we do not expect to be able to keep the Intrinsic Mass Transport Principle as is. If a sequence of finite graphs is FO-convergent to some modeling $\mathbf{L}$ then we require the following condition to hold, which we call *Finitary Mass Transport Principle* (FMTP):

For every measurable subsets of vertices A and B, if it holds $\deg_B(x) \geq a$ for every $x \in A$ and $\deg_A(y) \leq b$ for every $y \in B$ then $a\,\nu_{\mathbf{L}}(A) \leq b\,\nu_{\mathbf{L}}(B)$.

Note that in the case of modelings with bounded degrees, the Finitary Mass Transport Principle is equivalent to the Intrinsic Mass Transport Principle. Also note that the above equation holds necessarily when A and B are first-order definable, according to the convergence of the Stone pairings and the fact that the Finitary Mass Transport Principle obviously holds for finite graphs.

The second item of Theorem 1.9 will be quite difficult to establish and is the main result of this paper. In this later case, we obtain an inverse theorem:

THEOREM 1.11. *Every sequence of finite rooted colored trees with height at most C has a modeling* FO*-limit that is a rooted colored trees with height at most C, is weakly uniform, and satisfies the Finitary Mass Transport Principle.*

Conversely, every rooted colored tree modeling with height at most C that satisfies the Finitary Mass Transport Principle is the FO-limit of a sequence of finite rooted colored trees.

By Theorem 1.8, modeling FO-limit do not exist in general. However, we have a general representation of the limit of an FO-convergent sequence of λ-structures by means of a probability distribution on a compact Polish space S_λ defined from $\mathrm{FO}(\lambda)$ using Stone duality:

THEOREM 1.12. *Let λ be a fixed finite or countable signature. Then there exist two mappings $\mathbf{A} \mapsto \mu_{\mathbf{A}}$ and $\phi \mapsto K(\phi)$ such that*

- *$\mathbf{A} \mapsto \mu_{\mathbf{A}}$ is an injective mapping from the class of finite λ-structures to the space of regular probability measures on S_λ,*
- *$\phi \mapsto K(\phi)$ is a mapping from $\mathrm{FO}(\lambda)$ to the set of the clopen subsets of S_λ,*

such that for every finite λ-structure $\mathbf{A}$ and every first-order formula $\phi \in \mathrm{FO}(\lambda)$ it holds:

$$\langle \phi, \mathbf{A} \rangle = \int_{S_\lambda} 1_{K(\phi)} \, \mathrm{d}\mu_{\mathbf{A}}.$$

(To prevent risks of notational ambiguity, we shall use μ as root symbol for measures on Stone spaces and keep ν for measures on modelings.)

Consider an FO-convergent sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$. Then the pointwise convergence of $\langle\,\cdot\,, \mathbf{A}_n\rangle$ translates as a weak $*$-convergence of the measures $\mu_{\mathbf{A}_n}$ and we get:

THEOREM 1.13. *A sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of finite λ-structures is FO-convergent if and only if the sequence $(\mu_{\mathbf{A}_n})_{n\in\mathbb{N}}$ is weakly $*$-convergent.*

Moreover, if $\mu_{\mathbf{A}_n} \Rightarrow \mu$ then for every first-order formula $\phi \in \mathrm{FO}(\lambda)$ it holds:

$$\int_{S_\lambda} 1_{K(\phi)} \, \mathrm{d}\mu = \lim_{n\to\infty} \langle \phi, \mathbf{A}_n \rangle.$$

These last two Theorems are established in the next section as a warm up for our general theory.

CHAPTER 2

General Theory

2.1. Limits as Measures on Stone Spaces

In order to prove the representation theorems Theorem 1.12 and Theorem 1.13, we first need to prove a general representation for additive functions on Boolean algebras.

2.1.1. Representation of Additive Functions. Recall that a *Boolean algebra* $B = (B, \wedge, \vee, \neg, 0, 1)$ is an algebra with two binary operations $\vee$ and $\wedge$, a unary operation $\neg$ and two elements 0 and 1, such that $(B, \vee, \wedge)$ is a complemented distributive lattice with minimum 0 and maximum 1. The two-elements Boolean algebra is denoted $\mathbf{2}$.

To a Boolean algebra B is associated a topological space, denoted $S(B)$, whose points are the ultrafilters on B (or equivalently the homomorphisms $B \to \mathbf{2}$). The topology on $S(B)$ is generated by a sub-basis consisting of all sets

$$K_B(b) = \{x \in S(B) : b \in x\},$$

where $b \in B$. When the considered Boolean algebra will be clear from context we shall omit the subscript and write $K(b)$ instead of $K_B(b)$.

A topological space is a *Stone space* if it is Hausdorff, compact, and has a basis of clopen subsets. Boolean spaces and Stone spaces are equivalent as formalized by Stone representation theorem [74], which states (in the language of category theory) that there is a duality between the category of Boolean algebras (with homomorphisms) and the category of Stone spaces (with continuous functions). This justifies to call $S(B)$ the *Stone space* of the Boolean algebra B. The two contravariant functors defining this duality are denoted by S and Ω and defined as follows:

For every homomorphism $h : A \to B$ between two Boolean algebra, we define the map $S(h) : S(B) \to S(A)$ by $S(h)(g) = g \circ h$ (where points of $S(B)$ are identified with homomorphisms $g : B \to \mathbf{2}$). Then for every homomorphism $h : A \to B$, the map $S(h) : S(B) \to S(A)$ is a continuous function.

Conversely, for every continuous function $f : X \to Y$ between two Stone spaces, define the map $\Omega(f) : \Omega(Y) \to \Omega(X)$ by $\Omega(f)(U) = f^{-1}(U)$ (where elements of $\Omega(X)$ are identified with clopen sets of X). Then for every continuous function $f : X \to Y$, the map $\Omega(f) : \Omega(Y) \to \Omega(X)$ is a homomorphism of Boolean algebras.

We denote by $K = \Omega \circ S$ one of the two natural isomorphisms defined by the duality. Hence, for a Boolean algebra B, $K(B)$ is the set algebra $\{K_B(b) : b \in B\}$, and this algebra is isomorphic to B.

An ultrafilter of a Boolean algebra B can be considered as a finitely additive measure, for which every subset has either measure 0 or 1. Because of the equivalence of the notions of Boolean algebra and of set algebra, we define the space ba(B) as the space of all bounded additive functions $f : B \to \mathbb{R}$. Recall that a function $f : B \to \mathbb{R}$ is *additive* if for all $x, y \in B$ it holds

$$x \wedge y = 0 \quad \Longrightarrow \quad f(x \vee y) = f(x) + f(y).$$

The space ba(B) is a Banach space for the norm

$$\|f\|_{\mathrm{ba}(B)} = \sup_{x \in B} f(x) - \inf_{x \in B} f(x).$$

(Recall that the ba space of an algebra of sets Σ is the Banach space consisting of all bounded and finitely additive measures on Σ with the total variation norm.)

Let $V(B)$ be the normed vector space (of so-called *simple functions*) generated by the indicator functions of the clopen sets (equipped with supremum norm). The indicator function of clopen set $K(b)$ (for some $b \in B$) is denoted by $\mathbf{1}_{K(b)}$.

Lemma 2.1. *The space* ba(B) *is the topological dual of* $V(B)$

Proof. One can identify ba(B) with the space ba($K(B)$) of finitely additive measure defined on the set algebra $K(B)$. As a vector space, ba(B) $\approx$ ba($K(B)$) is then clearly the (algebraic) dual of the normed vector space $V(B)$.

The pairing of a function $f \in$ ba(B) and a vector $X = \sum_{i=1}^n a_i \mathbf{1}_{K(b_i)}$ is defined by

$$[f, X] = \sum_{i=1}^n a_i f(b_i).$$

That $[f, X]$ does not depend on a particular choice of a decomposition of X follows from the additivity of f. We include a short proof for completeness: Assume $\sum_i \alpha_i \mathbf{1}_{K(b_i)} = \sum_i \beta_i \mathbf{1}_{K(b_i)}$. As for every $b, b' \in B$ it holds $f(b) = f(b \wedge b') + f(b \wedge \neg b')$ and $\mathbf{1}_{K(b)} = \mathbf{1}_{K(b \wedge b')} + \mathbf{1}_{K(b \wedge \neg b')}$ we can express the two sums as $\sum_j \alpha'_j \mathbf{1}_{K(b'_j)} = \sum_j \beta'_j \mathbf{1}_{K(b'_j)}$ (where $b'_i \wedge b'_j = 0$ for every $i \neq j$), with $\sum_i \alpha_i f(b_i) = \sum_j \alpha'_j f(b'_j)$ and $\sum_i \beta_i f(b_i) = \sum_j \beta'_j f(b'_j)$. As $b'_i \wedge b'_j = 0$ for every $i \neq j$, for $x \in K(b'_j)$ it holds $\alpha'_j = X(x) = \beta'_j$. Hence $\alpha'_j = \beta'_j$ for every j. Thus $\sum_i \alpha_i f(b_i) = \sum_i \beta_i f(b_i)$.

Note that $X \mapsto [f, X]$ is indeed continuous. Thus ba(B) is the topological dual of $V(B)$. $\square$

Lemma 2.2. *The vector space* $V(B)$ *is dense in* $C(S(B))$ *(with the uniform norm).*

Proof. Let $f \in C(S(B))$ and let $\epsilon > 0$. For $z \in f(S(B))$ let U_z be the preimage by f of the open ball $B_{\epsilon/2}(z)$ of $\mathbb{R}$ centered in z. As f is continuous, U_z is a open set of $S(B)$. As $\{K(b) : b \in B\}$ is a basis of the topology of $S(B)$, U_z can be expressed as a union $\bigcup_{b \in \mathcal{F}(U_z)} K(b)$. It follows that $\bigcup_{z \in f(S(B))} \bigcup_{b \in \mathcal{F}(U_z)} K(b)$ is a covering of $S(B)$ by open sets. As $S(B)$ is compact, there exists a finite subset $\mathcal{F}$ of $\bigcup_{z \in f(S(B))} \mathcal{F}(U_z)$ that covers $S(B)$. Moreover, as for every $b, b' \in B$ it holds $K(b) \cap K(b') = K(b \wedge b')$ and $K(b) \setminus K(b') = K(b \wedge \neg b')$ it follows that we can assume that there exists a finite family $\mathcal{F}'$ such that $S(B)$ is covered by open sets $K(b)$ (for $b \in \mathcal{F}'$) and such that for every $b \in \mathcal{F}'$ there exists $b' \in \mathcal{F}$ such that $K(b) \subseteq K(b')$. In particular, it follows that for every $b \in \mathcal{F}'$, $f(K(b))$ is included in an open ball of radius $\epsilon/2$ of $\mathbb{R}$. For each $b \in \mathcal{F}'$ choose a point $x_b \in S(B)$ such that $b \in x_b$. Now define

$$\hat{f} = \sum_{b \in \mathcal{F}'} f(x_b) \mathbf{1}_{K(b)}$$

Let $x \in S(B)$. Then there exists $b \in \mathcal{F}'$ such that $x \in K(b)$. Thus

$$|f(x) - \hat{f}(x)| = |f(x) - f(x_b)| < \epsilon.$$

Hence $\|f - \hat{f}\|_\infty < \epsilon$. $\square$

LEMMA 2.3. *Let B be a Boolean algebra, let $\mathrm{ba}(B)$ be the Banach space of bounded additive real-valued functions equipped with the norm $\|f\| = \sup_{b \in B} f(b) - \inf_{b \in B} f(b)$, let $S(B)$ be the Stone space associated to B by Stone representation theorem, and let $\mathrm{rca}(S(B))$ be the Banach space of the regular countably additive measure on $S(B)$ equipped with the total variation norm.*

Then the mapping $C_K : \mathrm{rca}(S(B)) \to \mathrm{ba}(B)$ defined by $C_K(\mu) = \mu \circ K$ is an isometric isomorphism. In other words, C_K is defined by

$$C_K(\mu)(b) = \mu(\{x \in S(B) : \; b \in x\})$$

(considering that the points of $S(B)$ are the ultrafilters on B).

PROOF. According to Lemma 2.1, the Banach space $\mathrm{ba}(B)$ is the topological dual of $V(B)$ and as $V(B)$ is dense in $C(S(B))$ (according to Lemma 2.2) we deduce that $\mathrm{ba}(B)$ can be identified with the continuous dual of $C(S(B))$. By Riesz representation theorem, the topological dual of $C(S(B))$ is the space $\mathrm{rca}(S(B))$ of regular countably additive measures on $S(B)$. ¿From these observations follows the equivalence of $\mathrm{ba}(B)$ and $\mathrm{rca}(S(B))$.

This equivalence is easily made explicit, leading to the conclusion that the mapping $C_K : \mathrm{rca}(S(B)) \to \mathrm{ba}(B)$ defined by $C_K(\mu) = \mu \circ K$ is an isometric isomorphism. $\square$

Note also that, similarly, the restriction of C_K to the space $\Pr(S(B))$ of all (regular) probability measures on $S(B)$ is an isometric isomorphism of $\Pr(S(B))$ and the subset $\mathrm{ba}_1(B)$ of $\mathrm{ba}(B)$ of all positive additive functions f on B such that $f(1) = 1$.

Recall that given a measurable function $f : X \to Y$ (where X and Y are measurable spaces), the *pushforward* $f_*(\mu)$ of a measure μ on X is the measure on Y defined by $f_*(\mu)(A) = \mu(f^{-1}(A))$ (for every measurable set A of Y). Note that if f is a continuous function and if μ is a regular measure on X, then the pushforward measure $f_*(\mu)$ is a regular measure on Y. By similarity with the definition of $\Omega(f) : \Omega(Y) \to \Omega(X)$ (see above definition) we denote by $\Omega_*(f)$ the mapping from $\mathrm{rca}(X)$ to $\mathrm{rca}(Y)$ defined by $(\Omega_*(f))(\mu) = f_*(\mu)$.

All the functors defined above are consistent in the sense that if $h : A \to B$ is a homomorphism and $f \in \mathrm{ba}(B)$ then

$$\Omega_*(S(h))(\mu_f) \circ K_A = f \circ h = \tau_h(f).$$

A standard notion of convergence in $\mathrm{rca}(S(B))$ (as the continuous dual of $C(S(B))$) is the weak $*$-convergence: a sequence $(\mu_n)_{n \in \mathbb{N}}$ of measures is convergent if, for every $f \in C(S(B))$ the sequence $\int f(x) \, \mathrm{d}\mu_n(x)$ is convergent. Thanks to the density of $V(B)$ this convergence translates as pointwise convergence in $\mathrm{ba}(B)$ as follows: a sequence $(g_n)_{n \in \mathbb{N}}$ of functions in $\mathrm{ba}(B)$ is convergent if, for every $b \in B$ the sequence $(g_n(b))_{n \in \mathbb{N}}$ is convergent. As $\mathrm{rca}(S(B))$ is complete, so is $\mathrm{rca}(B)$. Moreover, it is easily checked that $\mathrm{ba}_1(B)$ is closed in $\mathrm{ba}(B)$.

In a more concise way, we can write, for a sequence $(f_n)_{n\in\mathbb{N}}$ of functions in $\mathrm{ba}(B)$ and for the corresponding sequence $(\mu_{f_n})_{n\in\mathbb{N}}$ of regular measures on $S(B)$:

$$f_n \to f \text{ pointwise} \qquad \Longleftrightarrow \qquad \mu_{f_n} \Rightarrow \mu_f.$$

We now apply this classical machinery to structures and models.

2.1.2. Basics of Model Theory and Lindenbaum-Tarski Algebras. We denote by $\mathcal{B}(\mathrm{FO}(\lambda))$ the equivalence classes of $\mathrm{FO}(\lambda)$ defined by logical equivalence. The (class of) unsatisfiable formulas (resp. of tautologies) will be designated by 0 (resp. 1). Then, $\mathcal{B}(\mathrm{FO}(\lambda))$ gets a natural structure of Boolean algebra (with minimum 0, maximum 1, infimum $\wedge$, supremum $\vee$, and complement $\neg$). This algbera is called the *Lindenbaum-Tarski algebra* of $\mathrm{FO}(\lambda)$. Notice that all the Boolean algebras $\mathrm{FO}(\lambda)$ for countable λ are isomorphic, as there exists only one countable atomless Boolean algebra up to isomorphism (see [41]).

For an integer $p \geq 1$, the fragment $\mathrm{FO}_p(\lambda)$ of $\mathrm{FO}(\lambda)$ contains first-order formulas ϕ such that $\mathrm{Fv}(\phi) \subseteq \{x_1, \dots, x_p\}$. The fragment $\mathrm{FO}_0(\lambda)$ of $\mathrm{FO}(\lambda)$ contains first-order formulas without free variables (that is *sentences*).

We check that the permutation group S_p on $[p]$ acts on $\mathrm{FO}_p(\lambda)$ by $\sigma \cdot \phi = \phi(x_{\sigma(1)}, \dots, x_{\sigma(p)})$ and that each permutation indeed define an automorphism of $\mathcal{B}(\mathrm{FO}_p(\lambda))$. Similarly, the group S_ω of permutation on $\mathbb{N}$ with finite support acts on $\mathrm{FO}(\lambda)$ and $\mathcal{B}(\mathrm{FO}(\lambda))$. Note that $\mathrm{FO}_0(\lambda) \subseteq \cdots \subseteq \mathrm{FO}_p(\lambda) \subseteq \mathrm{FO}_{p+1}(\lambda) \subseteq \cdots \subseteq \mathrm{FO}(\lambda)$. Conversely, let $\mathrm{rank}(\phi) = \max\{i : x_i \in \mathrm{Fv}(\phi)\}$. Then we have a natural projection $\pi_p : \mathrm{FO}(\lambda) \to \mathrm{FO}_p(\lambda)$ defined by

$$\pi_p(\phi) = \begin{cases} \phi & \text{if } \mathrm{rank}(\phi) \leq p \\ \exists x_{p+1}\, \exists x_{p+2}\, \dots\, \exists x_{\mathrm{rank}(\phi)}\, \phi & \text{otherwise} \end{cases}$$

An *elementary class* (or *axiomatizable class*) $\mathcal{C}$ of λ-structures is a class consisting of all λ-structures satisfying a fixed consistent first-order theory $T_{\mathcal{C}}$. Denoting by $\mathcal{I}_{T_{\mathcal{C}}}$ the ideal of all first-order formulas in $\mathcal{L}$ that are provably false from axioms in $T_{\mathcal{C}}$, The Lindenbaum-Tarski algebra $\mathcal{B}(\mathrm{FO}(\lambda), T_{\mathcal{C}})$ associated to the theory $T_{\mathcal{C}}$ of $\mathcal{C}$ is the quotient Boolean algebra $\mathcal{B}(\mathrm{FO}(\lambda), T_{\mathcal{C}}) = \mathcal{B}(\mathrm{FO}(\lambda))/\mathcal{I}_{T_{\mathcal{C}}}$. As a set, $\mathcal{B}(\mathrm{FO}(\lambda), T_{\mathcal{C}})$ is simply the quotient of $\mathrm{FO}(\lambda)$ by logical equivalence modulo $T_{\mathcal{C}}$.

As we consider countable languages, $T_{\mathcal{C}}$ is at most countable and it is easily checked that $S(\mathcal{B}(\mathrm{FO}(\lambda), T_{\mathcal{C}}))$ is homeomorphic to the compact subspace of $S(\mathcal{B}(\mathrm{FO}(\lambda))$ defined as $\{T \in S(\mathcal{B}(\mathrm{FO}(\lambda))) : T \supseteq T_{\mathcal{C}}\}$. Note that, for instance, $S(\mathcal{B}(\mathrm{FO}_0(\lambda), T_{\mathcal{C}}))$ is a clopen set of $S(\mathcal{B}(\mathrm{FO}_0(\lambda)))$ if and only if $\mathcal{C}$ is *finitely axiomatizable* (or a *basic* elementary class), that is if $T_{\mathcal{C}}$ can be chosen to be a single sentence. These explicit correspondences are particularly useful to our setting.

2.1.3. Stone Pairing Again. We add a few comments to Definition 1.6. Note first that this definition is consistent in the sense that for every modeling $\mathbf{A}$ and for every formula $\phi \in \mathrm{FO}(\lambda)$ with p free variables can be considered as a formula with $q \geq p$ free variables with $q - p$ unused variables, we have

$$\int_{A^q} 1_{\Omega_\phi(\mathbf{A})}(x)\, \mathrm{d}\nu_{\mathbf{A}}^q(x) = \int_{A^p} 1_{\Omega_\phi(\mathbf{A})}(x)\, \mathrm{d}\nu_{\mathbf{A}}^p(x).$$

It is immediate that for every formula ϕ it holds $\langle \neg\phi, \mathbf{A}\rangle = 1 - \langle \phi, \mathbf{A}\rangle$. Moreover, if $\phi_1, \dots, \phi_n$ are formulas, then by de Moivre's formula, it holds

$$\langle\bigvee_{i=1}^{n}\phi_i,\mathbf{A}\rangle=\sum_{k=1}^{n}(-1)^{k+1}\Bigl(\sum_{1\leq i_1<\cdots<i_k\leq n}\langle\bigwedge_{j=1}^{k}\phi_{i_j},\mathbf{A}\rangle\Bigr).$$

In particular, if $\phi_1,\dots,\phi_k$ are *mutually exclusive* (meaning that $\phi_i\wedge\phi_j=0$) then it holds

$$\langle\bigvee_{i=1}^{k}\phi_i,\mathbf{A}\rangle=\sum_{i=1}^{k}\langle\phi_i,\mathbf{A}\rangle.$$

It follows that for every fixed modeling $\mathbf{A}$, the mapping $\phi\mapsto\langle\phi,\mathbf{A}\rangle$ is additive (i.e. $\langle\,\cdot\,,\mathbf{A}\rangle\in\mathrm{ba}(\mathcal{B}(\mathrm{FO}(\lambda)))$):

$$\phi_1\wedge\phi_2=0\quad\Longrightarrow\quad\langle\phi_1\vee\phi_2,\mathbf{A}\rangle=\langle\phi_1,\mathbf{A}\rangle+\langle\phi_2,\mathbf{A}\rangle.$$

The Stone pairing is antimonotone:

Let $\phi,\psi\in\mathrm{FO}(\lambda)$. For every modeling $\mathbf{A}$ it holds

$$\phi\vdash\psi\quad\Longrightarrow\quad\langle\phi,G\rangle\geq\langle\psi,G\rangle.$$

However, even if ϕ and ψ are sentences and $\langle\phi,\,\cdot\,\rangle\geq\langle\psi,\,\cdot\,\rangle$ on finite λ-structures, this does not imply in general that $\phi\vdash\psi$: let θ be a sentence with only infinite models and let ϕ be a sentence with only finite models. On finite λ-structures it holds $\langle\phi\vee\theta,\,\cdot\,\rangle=\langle\phi,\,\cdot\,\rangle$ although $\phi\vee\theta\nvdash\phi$ (as witnessed by an infinite model of θ).

Nevertheless, inequalities between Stone pairing that are valid for finite λ-structures will of course still hold at the limit. For instance, for $\phi_1,\phi_2\in\mathrm{FO}_1(\lambda)$, for $\zeta\in\mathrm{FO}_2(\lambda)$, and for $a,b\in\mathbb{N}$ define the first-order sentence $B(a,b,\phi_1,\phi_2,\zeta)$ expressing that for every vertex x such that $\phi_1(x)$ holds there exist at least a vertices y such that $\phi_2(y)\wedge\zeta(x,y)$ holds and that for every vertex y such that $\phi_2(x)$ holds there exist at most b vertices x such that $\phi_1(x)\wedge\zeta(x,y)$ holds. Then it is easily checked that for every finite λ-structure $\mathbf{A}$ it holds

$$\mathbf{A}\models B(a,b,\phi_1,\phi_2,\zeta)\quad\Longrightarrow\quad a\langle\phi_1,\mathbf{A}\rangle\leq b\langle\phi_2,\mathbf{A}\rangle.$$

For example, if a finite directed graph is such that every arc connects a vertex with out-degree 2 to a vertex with in-degree 1, it is clear that the probability that a random vertex has out-degree 2 is half the probability that a random vertex has in-degree 1.

Now we come to important twist and the basic of our approach. The Stone pairing $\langle\,\cdot\,,\,\cdot\,\rangle$ can be considered from both sides: On the right side the functions of type $\langle\phi,\,\cdot\,\rangle$ are a generalization of the homomorphism density functions [11]:

$$t(F,G)=\frac{|\hom(F,G)|}{|G|^{|F|}}$$

(these functions correspond to $\langle\phi,G\rangle$ for Boolean conjunctive queries ϕ and a graph G). Also the density function used in [8] to measure the probability that the ball of radius r rooted at a random vertex as a given isomorphism type may be expressed as a function $\langle\phi,\,\cdot\,\rangle$. We follow here, in a sense, a dual approach: from the left side we consider for fixed $\mathbf{A}$ the function $\langle\,\cdot\,,\mathbf{A}\rangle$, which is an additive function on $\mathcal{B}(\mathrm{FO}(\lambda))$ with the following properties:

- $\langle\,\cdot\,,\mathbf{A}\rangle\geq 0$ and $\langle 1,\mathbf{A}\rangle=1$;
- $\langle\sigma\cdot\phi,\mathbf{A}\rangle=\langle\phi,\rangle$ for every $\sigma\in S_\omega$;

- if $\mathrm{Fv}(\phi) \cap \mathrm{Fv}(\psi) = \emptyset$, then $\langle \phi \wedge \psi, \mathbf{A} \rangle = \langle \phi, \mathbf{A} \rangle \; \langle \psi, \mathbf{A} \rangle$.

Thus $\langle \cdot \, , \mathbf{A} \rangle$ is, for a given $\mathbf{A}$, an operator on the class of first-order formulas.

We now can apply Lemma 2.3 to derive a representation by means of a regular measure on a Stone space. The fine structure and interplay of additive functions, Boolean functions, and dual spaces can be used effectively if we consider finite λ-structures as probability spaces as we did when we considered finite λ-structures as a particular case of Borel models.

The following two theorems generalize Theorems 1.12 and 1.13 mentioned in Section 1.1.

THEOREM 2.4. *Let λ be a signature, let $\mathcal{B}(\mathrm{FO}(\lambda))$ be the Lindenbaum-Tarski algebra of $\mathrm{FO}(\lambda)$, let $S(\mathcal{B}(\mathrm{FO}(\lambda)))$ be the associated Stone space, and let $\mathrm{rca}(S(\mathcal{B}(\mathrm{FO}(\lambda))))$ be the Banach space of the regular countably additive measure on $S(\mathcal{B}(\mathrm{FO}(\lambda)))$. Then:*

(1) There is a mapping from the class of λ-modeling to $\mathrm{rca}(S(\mathcal{B}(\mathrm{FO}(\lambda))))$, which maps a modeling $\mathbf{A}$ to the unique regular measure $\mu_{\mathbf{A}}$ such that for every $\phi \in \mathrm{FO}(\lambda)$ it holds

$$\langle \phi, \mathbf{A} \rangle = \int_{S(\mathcal{B}(\mathrm{FO}(\lambda)))} \mathbf{1}_{K(\phi)} \, \mathrm{d}\mu_{\mathbf{A}},$$

where $\mathbf{1}_{K(\phi)}$ is the indicator function of $K(\phi)$ in $S(\mathcal{B}(\mathrm{FO}(\lambda)))$. Moreover, this mapping is injective of finite λ-structures.

(2) A sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of finite λ-structures is FO-convergent if and only if the sequence $(\mu_{\mathbf{A}_n})_{n\in\mathbb{N}}$ is weakly converging in $\mathrm{rca}(S(\mathcal{B}(\mathrm{FO}(\lambda))))$;

(3) If $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is an FO-convergent sequence of finite λ-structures then the weak limit μ of $(\mu_{\mathbf{A}_n})_{n\in\mathbb{N}}$ is such that for every $\phi \in \mathrm{FO}(\lambda)$ it holds

$$\lim_{n\to\infty} \langle \phi, \mathbf{A}_n \rangle = \int_{S(\mathcal{B}(\mathrm{FO}(\lambda)))} \mathbf{1}_{K(\phi)} \, \mathrm{d}\mu.$$

PROOF. The proof follows from Lemma 2.3, considering the additive functions $\langle \cdot \, , \mathbf{A} \rangle$.

Let $\mathbf{A}$ be a finite λ-structure. As $\mu_{\mathbf{A}}$ allows to recover the complete theory of $\mathbf{A}$ and as $\mathbf{A}$ is finite, the mapping $\mathbf{A} \mapsto \mu_{\mathbf{A}}$ is injective. $\square$

It is important to consider fragments of $\mathrm{FO}(\lambda)$ to define a weaker notion of convergence. This allows us to capture limits of dense graphs too.

DEFINITION 2.5 (X-convergence). Let X be a fragment of $\mathrm{FO}(\lambda)$. A sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of finite λ-structures is *X-convergent* if $\langle \phi, \mathbf{A}_n \rangle$ is convergent for every $\phi \in X$.

In the particular case that X is a Boolean sub algebra of $\mathcal{B}(\mathrm{FO}(\lambda))$ we can apply all above methods and in this context we can extend Theorem 2.4.

THEOREM 2.6. *Let λ be a signature, and let X be a fragment of $\mathrm{FO}(\lambda)$ defining a Boolean algebra $\mathcal{B}(X) \subseteq \mathcal{B}(\mathrm{FO}(\lambda))$. Let $S(\mathcal{B}(X))$ be the associated Stone space, and let $\mathrm{rca}(S(\mathcal{B}(X)))$ be the Banach space of the regular countably additive measure on $S(\mathcal{B}(X))$. Then:*

(1) *The canonical injection* $\iota^X : \mathcal{B}(X) \to \mathcal{B}(\mathrm{FO}(\lambda))$ *defines by duality a continuous projection* $p^X : S(\mathcal{B}(\mathrm{FO}(\lambda))) \to S(\mathcal{B}(X))$; *The pushforward* $p^X_* \mu_{\mathbf{A}}$ *of the measure* $\mu_{\mathbf{A}}$ *associated to a modeling* $\mathbf{A}$ *(see Theorem 2.4) is the unique regular measure on* $S(\mathcal{B}(X))$ *such that:*

$$\langle \phi, \mathbf{A} \rangle = \int_{S(\mathcal{B}(X))} \mathbf{1}_{K(\phi)} \, \mathrm{d}p^X_* \mu_{\mathbf{A}},$$

where $\mathbf{1}_{K(\phi)}$ *is the indicator function of* $K(\phi)$ *in* $S(\mathcal{B}(X))$.

(2) *A sequence* $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *of finite* λ*-structures is* X*-convergent if and only if the sequence* $(p^X_* \mu_{\mathbf{A}_n})_{n\in\mathbb{N}}$ *is weakly converging in* $\mathrm{rca}(S(\mathcal{B}(X)))$;

(3) *If* $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *is an* X*-convergent sequence of finite* λ*-structures then the weak limit* μ *of* $(p^X_* \mu_{\mathbf{A}_n})_{n\in\mathbb{N}}$ *is such that for every* $\phi \in X$ *it holds*

$$\lim_{n\to\infty} \langle \phi, \mathbf{A}_n \rangle = \int_{S(\mathcal{B}(X))} \mathbf{1}_{K(\phi)} \, \mathrm{d}\mu.$$

Proof. If X is closed under conjunction, disjunction and negation, thus defining a Boolean algebra $\mathcal{B}(X)$, then the inclusion of X in $\mathrm{FO}(\lambda)$ translates as a canonical injection ι from $\mathcal{B}(X)$ to $\mathcal{B}(\mathrm{FO}(\lambda))$. By Stone duality, the injection ι corresponds to a continuous projection $p : S(\mathcal{B}(\mathrm{FO}(\lambda))) \to S(\mathcal{B}(X))$. As every measurable function, this continuous projection also transports measures by pushforward: the projection p transfers the measure μ on $S(\mathcal{B}(\mathrm{FO}(\lambda)))$ to $S(\mathcal{B}(X))$ as the pushforward measure $p_*\mu$ defined by the identity $p_*\mu(Y) = \mu(p^{-1}(Y))$, which holds for every measurable subset Y of $S(\mathcal{B}(X))$.

The proof follows from Lemma 2.3, considering the additive functions $\langle \cdot, \mathbf{A} \rangle$. □

We can also consider a notion of convergence restricted to λ-structures satisfying a fixed axiom.

Theorem 2.7. *Let* λ *be a signature, and let* X *be a fragment of* $\mathrm{FO}(\lambda)$ *defining a Boolean algebra* $\mathcal{B}(X) \subseteq \mathcal{B}(\mathrm{FO}(\lambda))$. *Let* $S(\mathcal{B}(X))$ *be the associated Stone space, and let* $\mathrm{rca}(S(\mathcal{B}(X)))$ *be the Banach space of the regular countably additive measure on* $S(\mathcal{B}(X))$.

Let $\mathcal{C}$ *be a basic elementary class defined by a single axiom* $\Psi \in X \cap \mathrm{FO}_0$, *and let* $\mathcal{I}_\Psi$ *be the principal ideal of* $\mathcal{B}(X)$ *generated by* $\neg\Psi$.

Then:

(1) *The Boolean algebra obtained by taking the quotient of* X *equivalence modulo* Ψ *is the quotient Boolean algebra* $\mathcal{B}(X, \Psi) = \mathcal{B}(X)/\mathcal{I}_\Psi$. *Then* $S(\mathcal{B}(X, \Psi))$ *is homeomorphic to the clopen subspace* $K(\Psi)$ *of* $S(\mathcal{B}(X))$.

If $\mathbf{A} \in \mathcal{C}$ *is a finite* λ*-structure then the support of the measure* $p^X_* \mu_{\mathbf{A}}$ *associated to* $\mathbf{A}$ *(see Theorem 2.6) is included in* $K(\Psi)$ *and for every* $\phi \in X$ *it holds*

$$\langle \phi, \mathbf{A} \rangle = \int_{K(\Psi)} \mathbf{1}_{K(\phi)} \, \mathrm{d}p^X_* \mu_{\mathbf{A}}.$$

(2) *A sequence* $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *of finite* λ*-structures of* $\mathcal{C}$ *is* X*-convergent if and only if the sequence* $(p^X_* \mu_{\mathbf{A}_n})_{n\in\mathbb{N}}$ *is weakly converging in* $\mathrm{rca}(S(\mathcal{B}(X, \Psi)))$;

(3) If $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is an X-convergent sequence of finite λ-structures in $\mathcal{C}$ then the weak limit μ of $(p_*^X\mu_{\mathbf{A}_n})_{n\in\mathbb{N}}$ is such that for every $\phi\in X$ it holds

$$\lim_{n\to\infty}\langle\phi,\mathbf{A}_n\rangle=\int_{K(\Psi)}\mathbf{1}_{K(\phi)}\,\mathrm{d}\mu.$$

Proof. The quotient algebra $\mathcal{B}(X,\Psi)=\mathcal{B}(X)/\mathcal{I}_\Psi$ is isomorphic to the sub-Boolean algebra $\mathcal{B}'$ of $\mathcal{B}$ of all (equivalence classes of) formulas $\phi\wedge\Psi$ for $\phi\in X$. To this isomorphism corresponds by duality the identification of $S(\mathcal{B}(X,\Psi))$ with the clopen subspace $K(\Psi)$ of $S(\mathcal{B}(X))$. $\square$

The situation expressed by these theorems is summarized in the following diagram.

$$\begin{tikzcd}
\mathcal{B}(\mathrm{FO}(\lambda)) \arrow[d, leftrightarrow] & \mathcal{B}(X) \arrow[l, "\text{canonical injection}"'] \arrow[d, leftrightarrow] & \mathcal{B}' \arrow[l, "\text{inclusion}"'] \arrow[r, leftrightarrow, "\textit{isomorphism}"] \arrow[d, leftrightarrow] & \mathcal{B}(X,\Psi) \arrow[d, leftrightarrow] \\
S(\mathcal{B}(\mathrm{FO}(\lambda))) \arrow[r, "\text{projection } p^X"] & S(\mathcal{B}(X)) & K(\Psi) \arrow[l, "\text{inclusion}"'] \arrow[r, leftrightarrow, "\text{homeomorphism}"] & S(\mathcal{B}(X,\Psi))
\end{tikzcd}$$

$$\mu \xrightarrow{\text{pushforward}} p_*^X\mu \xrightarrow{\text{restriction}} p_*^X\mu$$

The essence of our approach is that we follow a dual path: we view a graph G as an operator on first-order formulas through Stone pairing $\langle\,\cdot\,,G\rangle$.

2.1.4. Limit of Measures Associated to Finite Structures. We consider a signature λ and fragment FO_p of $\mathrm{FO}(\lambda)$. Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be an X-convergent sequence of λ-structures, let $\mu_{\mathbf{A}_n}$ be the measure on $S(\mathcal{B}(X))$ associated to $\mathbf{A}_n$, and let μ be the weak limit of $\mu_{\mathbf{A}_n}$.

Fact 2.8. As we consider countable languages only, $S(\mathcal{B}(\mathrm{FO}_p))$ is a Radon space and thus for every (Borel) probability measure μ on $S(\mathcal{B}(\mathrm{FO}_p))$, any measurable set outside the support of μ has zero μ-measure.

Definition 2.9. Let π be the natural projection $S(\mathcal{B}(\mathrm{FO}_p))\to S(\mathcal{B}(\mathrm{FO}_0))$.

A measure μ on $S(\mathcal{B}(\mathrm{FO}_p))$ is *pure* if $|\pi(\mathrm{Supp}(\mu))|=1$. The unique element T of $\pi(\mathrm{Supp}(\mu))$ is then called the *complete theory* of μ.

Remark 2.10. Every measure μ that is the weak limit of some sequence of measures associated to finite structures is pure and its complete theory has the finite model property.

Definition 2.11. For $T\in S(\mathcal{B}(\mathrm{FO}_p),\psi,\phi\in\mathrm{FO}_p$, and $\beta\in\mathrm{FO}_{2p}$ define

$$\deg_\psi^{\beta+}(T)=\begin{cases}k & \text{if } T\ni(\exists^{=k}(y_1,\dots,y_p)\,\beta(x_1,\dots,x_p,y_1,\dots,y_p)\wedge\psi(y_1,\dots,y_p))\\ \infty & \text{otherwise.}\end{cases}$$

$$\deg_\phi^{\beta-}(T)=\begin{cases}k & \text{if } T\ni(\exists^{=k}(x_1,\dots,x_p)\,\phi(x_1,\dots,x_p)\wedge\beta(x_1,\dots,x_p,y_1,\dots,y_p))\\ \infty & \text{otherwise.}\end{cases}$$

If μ is a measure associated to a finite structure then for every $\phi, \psi \in \mathrm{FO}_p$ it holds
$$\int_{K(\phi)} \deg_{\psi}^{\beta^+}(T)\,\mathrm{d}\mu(T) = \int_{K(\psi)} \deg_{\phi}^{\beta-}(T)\,\mathrm{d}\mu(T).$$

Hence for every measure μ that is the weak limit of some sequence of measures associated to finite structures the following property holds:

General Finitary Mass Transport Principle (GFMTP)
For every $\phi, \psi \in \mathrm{FO}_p$, every $\beta \in \mathrm{FO}_{2p}$, and every integers a, b that are such that
$$\begin{cases} \forall T \in K(\phi) & \deg_{\psi}^{\beta^+}(T) \geq a \\ \forall T \in K(\psi) & \deg_{\phi}^{\beta-}(T) \leq b \end{cases}$$
it holds
$$a\,\mu(K(\phi)) \leq b\,\mu(K(\psi)).$$

Of course, similar statement holds as well for the projection of μ on $S(\mathcal{B}(\mathrm{FO}_q))$ for $q < p$. In the case of digraphs, when $p = 1$ and $\beta(x_1, x_2)$ is existence of an arc from x_1 to x_2, we shall note $\deg_{\psi}^{+}$ and $\deg_{\phi}^{-}$ instead of $\deg_{\psi}^{\beta+}$ and $\deg_{\phi}^{\beta-}$. (In the case of graphs, we have $\deg_{\psi}^{+} = \deg_{\psi}^{-} = \deg_{\psi}$.) Thus it holds

Finitary Mass Transport Principle (FMTP)
For every $\phi, \psi \in \mathrm{FO}_1$, and every integers a, b that are such that
$$\begin{cases} \forall T \in K(\phi) & \deg_{\psi}^{+}(T) \geq a \\ \forall T \in K(\psi) & \deg_{\phi}^{-}(T) \leq b \end{cases}$$
it holds
$$a\,\mu(K(\phi)) \leq b\,\mu(K(\psi)).$$

GFMTP and FMTP will play a key role in the analysis of modeling limits.

2.2. Convergence, Old and New

As we have seen above, there are many ways how to say that a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of finite λ-structures is convergent. As we considered λ-structures defined with a countable signature λ, the Boolean algebra $\mathcal{B}(\mathrm{FO}(\lambda))$ is countable. It follows that the Stone space $S(\mathcal{B}(\mathrm{FO}(\lambda)))$ is a Polish space thus (with the Borel σ-algebra) it is a standard Borel space. Hence every probability distribution turns $S(\mathcal{B}(\mathrm{FO}(\lambda)))$ into a standard probability space. However, the fine structure of $S(\mathcal{B}(\mathrm{FO}(\lambda)))$ is complex and we have no simple description of this space.

FO-convergence is of course the most restrictive notion of convergence and it seems (at least on the first glance) that this is perhaps too much to ask, as we may encounter many particular difficulties and specific cases. But we shall exhibit later classes for which FO-convergence is captured — for special basic elementary classes of structures — by X-convergence for a small fragment X of FO.

At this time it is natural to ask whether one can consider fragments that are not sub-Boolean algebras of $\mathrm{FO}(\mathcal{L})$ and still have a description of the limit of a converging sequence as a probability measure on a nice measurable space. There is obviously a case where this is possible: when the convergence of $\langle \phi, \mathbf{A}_n \rangle$ for every ϕ in a fragment X implies the convergence of $\langle \psi, \mathbf{A}_n \rangle$ for every ψ in the minimum

Boolean algebra containing X. We prove now that this is for instance the case when X is a fragment closed under conjunction.

We shall need the following preliminary lemma:

Lemma 2.12. *Let $X \subseteq B$ be closed by $\wedge$ and such that X* generates *B (i.e. such that $B[X] = B$).*

Then $\{\mathbf{1}_b : \ b \in X\} \cup \{\mathbf{1}\}$ (where $\mathbf{1}$ is the constant function with value 1) includes a basis of the vector space $V(B)$ generated by the whole set $\{\mathbf{1}_b : \ b \in B\}$.

Proof. Let $b \in B$. As X generates B there exist $b_1, \dots, b_k \in X$ and a Boolean function F such that $b = F(b_1, \dots, b_k)$. As $\mathbf{1}_{x \wedge y} = \mathbf{1}_x \mathbf{1}_y$ and $\mathbf{1}_{\neg x} = \mathbf{1} - \mathbf{1}_x$ there exists a polynomial P_F such that $\mathbf{1}_b = P_F(\mathbf{1}_{b_1}, \dots, \mathbf{1}_{b_k})$. For $I \subseteq [k]$, the monomial $\prod_{i \in I} \mathbf{1}_{b_i}$ rewrites as $\mathbf{1}_{b_I}$ where $b_I = \bigwedge_{i \in I} b_i$. It follows that $\mathbf{1}_b$ is a linear combination of the functions $\mathbf{1}_{b_I}$ ($I \subseteq [k]$) which belong to X if $I \neq \emptyset$ (as X is closed under $\wedge$ operation) and equal $\mathbf{1}$, otherwise. □

Proposition 2.1. *Let X be a fragment of* $\mathrm{FO}(\lambda)$ *closed under (finite) conjunction — thus defining a* meet semilattice *of $\mathcal{B}(\mathrm{FO}(\lambda))$ — and let $\mathcal{B}(X)$ be the sub-Boolean algebra of $\mathcal{B}(\mathrm{FO}(\lambda))$ generated by X. Let $\overline{X}$ be the fragment of* $\mathrm{FO}(\lambda)$ *consisting of all formulas with equivalence class in $\mathcal{B}(X)$.*

Then X-convergence is equivalent to $\overline{X}$-convergence.

Proof. Let $\Psi \in \overline{X}$. According to Lemma 2.12, there exist $\phi_1, \dots, \phi_k \in X$ and $\alpha_0, \alpha_1, \dots, \alpha_k \in \mathbb{R}$ such that

$$\mathbf{1}_\Psi = \alpha_0 \mathbf{1} + \sum_{i=1}^k \alpha_i \mathbf{1}_{\phi_i}.$$

Let $\mathbf{A}$ be a λ-structure, let $\Omega = S(\mathcal{B}(X))$ and let $\mu_{\mathbf{A}} \in \mathrm{rca}(\Omega)$ be the associated measure. Then

$$\langle \Psi, \mathbf{A} \rangle = \int_\Omega \mathbf{1}_\Psi \,\mathrm{d}\mu_{\mathbf{A}} = \int_\Omega \big(\alpha_0 \mathbf{1} + \sum_{i=1}^k \alpha_i \mathbf{1}_{\phi_i}\big) \,\mathrm{d}\mu_G = \alpha_0 + \sum_{i=1}^k \alpha_i \langle \phi_i, \mathbf{A} \rangle.$$

It follows that if $(\mathbf{A}_n)_{n \in \mathbb{N}}$ is an X-convergent sequence, the sequence $(\langle \psi, \mathbf{A}_n \rangle)_{n \in \mathbb{N}}$ converges for every $\psi \in \overline{X}$, that is $(\mathbf{A}_n)_{n \in \mathbb{N}}$ is $\overline{X}$-convergent. □

Now we demonstrate the expressive power of X-convergence by relating it to the main types of convergence of graphs studied previously:

(1) the notion of *dense graph limit* [12, 54];
(2) the notion of *bounded degree graph limit* [8, 5];
(3) the notion of *elementary limit* derived from two important results in first-order logic, namely Gödel's completeness theorem and the compactness theorem.

These standard notions of graph limits, which have inspired this work, correspond to special fragments of $\mathrm{FO}(\lambda)$, where γ is the signature of graphs. In the remaining of this section, we shall only consider undirected graphs, thus we shall omit to precise their signature in the notations as well as the axiom defining the basic elementary class of undirected graphs.

2.2.1. L-convergence and QF**-convergence.** Recall that a sequence $(G_n)_{n\in\mathbb{N}}$ of graphs is *L-convergent* if

$$t(F, G_n) = \frac{\hom(F, G_n)}{|G_n|^{|F|}}$$

converges for every fixed (connected) graph F, where $\hom(F, G)$ denotes the number of homomorphisms of F to G [54, 13, 14].

It is a classical observation that homomorphisms between finite structures can be expressed by Boolean conjunctive queries [16]. We denote by HOM the fragment of FO consisting of formulas formed by conjunction of atoms. For instance, the formula

$$(x_1 \sim x_2) \wedge (x_2 \sim x_3) \wedge (x_3 \sim x_4) \wedge (x_4 \sim x_5) \wedge (x_5 \sim x_1)$$

belongs to HOM and it expresses that $(x_1, x_2, x_3, x_4, x_5)$ form a homomorphic image of C_5. Generally, to a finite graph F we associate the canonical formula $\phi_F \in$ HOM defined by

$$\phi_F := \bigwedge_{ij\in E(F)} (x_i \sim x_j).$$

Then, for every graph G it holds

$$\langle \phi_F, G\rangle = \frac{\hom(F, G)}{|G|^{|F|}} = t(F, G).$$

Thus L-convergence is equivalent to HOM-convergence. According to Proposition 2.1, HOM-convergence is equivalent to $\overline{\text{HOM}}$-convergent. It is easy to see that $\overline{\text{HOM}}$ is the fragment QF^- of quantifier free formulas that do not use equality. We prove now that HOM-convergence is actually equivalent to QF-convergence, where QF is the fragment of all quantifier free formulas. Note that QF is a proper fragment of FO^{local}.

THEOREM 2.13. *Let (G_n) be a sequence of finite graphs such that $\lim_{n\to\infty} |G_n| = \infty$.*

Then the following conditions are equivalent:

1. *the sequence (G_n) is L-convergent;*
2. *the sequence (G_n) is* QF^-*-convergent;*
3. *the sequence (G_n) is* QF*-convergent;*

PROOF. As L-convergence is equivalent to HOM-convergence and as HOM $\subset$ $\text{QF}^- \subset$ QF, it is sufficient to prove that L-convergence implies QF-convergence.

Assume (G_n) is L-convergent. The inclusion-exclusion principle implies that for every finite graph F the density of induced subgraphs isomorphic to F converges too. Define

$$\text{dens}(F, G_n) = \frac{(\#F \subseteq_i G_n)}{|G_n|^{|F|}}.$$

Then $\text{dens}(F, G_n)$ is a converging sequence for each F.

Let θ be a quantifier-free formula with $\text{Fv}(\theta) \subseteq [p]$. We first consider all possible cases of equalities between the free variables. For a partition $\mathcal{P} = (I_1, \ldots, I_k)$ of

$[p]$, we define $|\mathcal{P}| = k$ and $s_{\mathcal{P}}(i) = \min I_i$ (for $1 \leq i \leq |\mathcal{P}|$). Consider the formula

$$\zeta_{\mathcal{P}} := \bigwedge_{i=1}^{|\mathcal{P}|}\Big(\bigwedge_{j\in I_i}(x_j = x_{s_{\mathcal{P}}(i)}) \ \wedge \ \bigwedge_{j=i+1}^{|\mathcal{P}|}(x_{s_{\mathcal{P}}(j)} \neq x_{s_{\mathcal{P}}(i)})\Big).$$

Then θ is logically equivalent to

$$(\bigwedge_{i\neq j}(x_i \neq x_j) \wedge \theta) \ \vee \bigvee_{\mathcal{P}:|\mathcal{P}|<p} \zeta_{\mathcal{P}} \wedge \theta_{\mathcal{P}}(x_{s_{\mathcal{P}}(1)}, \dots, x_{s_{\mathcal{P}}(|\mathcal{P}|)}).$$

Note that all the formulas in the disjunction are mutually exclusive. Also $\bigwedge_{i\neq j}(x_i \neq x_j) \wedge \theta$ may be expressed as a disjunction of mutually exclusive terms:

$$\bigwedge_{i\neq j}(x_i \neq x_j) \wedge \theta = \bigvee_{F\in\mathcal{F}} \theta'_F,$$

where $\mathcal{F}$ is a finite family of finite graphs F and where $G \models \theta'_F(v_1, \dots, v_p)$ if and only if the mapping $i \mapsto v_i$ is an isomorphism from F to $G[v_1, \dots, v_p]$.

It follows that for every graph G it holds:

$$\begin{aligned}
\langle \theta, G\rangle &= \sum_{F\in\mathcal{F}} \langle \theta'_F, G\rangle + \sum_{\mathcal{P}:|\mathcal{P}|<p} \langle \zeta_{\mathcal{P}} \wedge \theta_{\mathcal{P}}(x_{s_{\mathcal{P}}(1)}, \dots, x_{s_{\mathcal{P}}(|\mathcal{P}|)}), G\rangle \\
&= \sum_{F\in\mathcal{F}} \langle \theta'_F, G\rangle + \sum_{\mathcal{P}:|\mathcal{P}|<p} |G|^{|\mathcal{P}|-p}\langle \theta_{\mathcal{P}}, G\rangle \\
&= \sum_{F\in\mathcal{F}} \frac{1}{p!} \sum_{\sigma\in\mathfrak{S}_p} \frac{|\{(v_1, \dots, v_p) : \ G \models \theta'_F(v_{\sigma(1)}, \dots, v_{\sigma(p)})\}|}{|G|^p} + O(|G|^{-1}) \\
&= \sum_{F\in\mathcal{F}} \frac{\mathrm{Aut}(F)}{p!} \operatorname{dens}(F, G) + O(|G|^{-1}).
\end{aligned}$$

Thus $\langle \theta, G_n\rangle$ converge for every quantifier-free formula θ. Hence (G_n) is QF-convergent. $\square$

Notice that the condition that $\lim_{n\to\infty} |G_n|$ is necessary as witnessed by the sequence (G_n) where G_n is K_1 if n is odd and $2K_1$ if n is even. The sequence is obviously L-convergent, but not QF convergent as witnessed by the formula $\phi(x, y) : x \neq y$, which has density 0 in K_1 and 1/2 in K_2.

Remark 2.14. The Stone space of the fragment QF^- has a simple description. Indeed, a homomorphism $h : \mathcal{B}(\mathrm{QF}^-) \to \mathbf{2}$ is determined by its values on the formulas $x_i \sim x_j$ and any mapping from this subset of formulas to $\mathbf{2}$ extends (in a unique way) to a homomorphism of $\mathcal{B}(\mathrm{QF}^-)$ to $\mathbf{2}$. Thus the points of $S(\mathcal{B}(\mathrm{QF}^-))$ can be identified with the mappings from $\binom{\mathbb{N}}{2}$ to $\{0, 1\}$ that is to the graphs on $\mathbb{N}$. Hence the considered measures μ are probability measures of graphs on $\mathbb{N}$ that have the property that they are invariant under the natural action of S_ω on $\mathbb{N}$. Such random graphs on $\mathbb{N}$ are called *infinite exchangeable random graphs*. For more on infinite exchangeable random graphs and graph limits, see e.g. [6, 21].

2.2.2. BS-convergence and $\mathrm{FO}^{\mathrm{local}}$-convergence. The class of graphs with maximum degree at most D (for some integer D) received much attention. Specifically, the notion of *local weak convergence* of bounded degree graphs was introduced in [8], which is called here *BS-convergence*:

A *rooted graph* is a pair (G, o), where $o \in V(G)$. An *isomorphism* of rooted graph $\phi : (G, o) \to (G', o')$ is an isomorphism of the underlying graphs which satisfies $\phi(o) = o'$. Let $D \in \mathbb{N}$. Let $\mathcal{G}_D$ denote the collection of all isomorphism classes of connected rooted graphs with maximal degree at most D. For the sake of simplicity, we denote elements of $\mathcal{G}_D$ simply as graphs. For $(G, o) \in \mathcal{G}_D$ and $r \geq 0$ let $B_G(o, r)$ denote the subgraph of G spanned by the vertices at distance at most r from o. If $(G, o), (G', o') \in \mathcal{G}_D$ and r is the largest integer such that $(B_G(o, r), o)$ is rooted-graph isomorphic to $(B_{G'}(o', r), o')$, then set $\rho((G, o), (G', o')) = 1/r$, say. Also take $\rho((G, o), (G, o)) = 0$. Then ρ is metric on $\mathcal{G}_D$. Let $\mathfrak{M}_D$ denote the space of all probability measures on $\mathcal{G}_D$ that are measurable with respect to the Borel σ-field of ρ. Then $\mathfrak{M}_D$ is endowed with the topology of weak convergence, and is compact in this topology.

A sequence $(G_n)_{n \in \mathbb{N}}$ of finite connected graphs with maximum degree at most D is *BS-convergent* if, for every integer r and every rooted connected graph (F, o) with maximum degree at most D the following limit exists:

$$\lim_{n \to \infty} \frac{|\{v : B_{G_n}(v, r) \cong (F, o)\}|}{|G_n|}.$$

This notion of limits leads to the definition of a limit object as a probability measure on $\mathcal{G}_D$ [8].

To relate BS-convergence to X-convergence, we shall consider the fragment of local formulas:

Let $r \in \mathbb{N}$. A formula $\phi \in \mathrm{FO}_p$ is *r-local* if, for every graph G and every $v_1, \dots, v_p \in G^p$ it holds

$$G \models \phi(v_1, \dots, v_p) \quad \iff \quad G[N_r(v_1, \dots, v_p)] \models \phi(v_1, \dots, v_p),$$

where $G[N_r(v_1, \dots, v_p)]$ denotes the subgraph of G induced by all the vertices at (graph) distance at most r from one of $v_1, \dots, v_p$ in G.

A formula ϕ is *local* if it is r-local for some $r \in \mathbb{N}$; the fragment $\mathrm{FO}^{\mathrm{local}}$ is the set of all local formulas in FO. Notice that if ϕ_1 and ϕ_2 are local formulas, so are $\phi_1 \wedge \phi_2$, $\phi_1 \vee \phi_2$ and $\neg\phi_1$. It follows that the quotient of $\mathrm{FO}^{\mathrm{local}}$ by the relation of logical equivalence defines a sub-Boolean algebra $\mathcal{B}(\mathrm{FO}^{\mathrm{local}})$ of $\mathcal{B}(\mathrm{FO})$. For $p \in \mathbb{N}$ we further define $\mathrm{FO}_p^{\mathrm{local}} = \mathrm{FO}^{\mathrm{local}} \cap \mathrm{FO}_p$.

THEOREM 2.15. *Let (G_n) be a sequence of finite graphs with maximum degree d, with $\lim_{n \to \infty} |G_n| = \infty$.*

Then the following properties are equivalent:

(1) the sequence $(G_n)_{n \in \mathbb{N}}$ is BS-convergent;
(2) the sequence $(G_n)_{n \in \mathbb{N}}$ is $\mathrm{FO}_1^{\mathrm{local}}$-convergent;
(3) the sequence $(G_n)_{n \in \mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent.

PROOF. If $(G_n)_{n \in \mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent, it is $\mathrm{FO}_1^{\mathrm{local}}$-convergent;

If $(G_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}_1^{\mathrm{local}}$-convergent then it is BS-convergent as for any finite rooted graph (F,o), testing whether the the ball of radius r centered at a vertex x is isomorphic to (F,o) can be formulated by a local first order formula.

Assume $(G_n)_{n\in\mathbb{N}}$ is BS-convergent. As we consider graphs with maximum degree d, there are only finitely many isomorphism types for the balls of radius r centered at a vertex. It follows that any local formula $\xi(x)$ with a single variable can be expressed as the conjunction of a finite number of (mutually exclusive) formulas $\xi_{(F,o)}(x)$, which in turn correspond to subgraph testing. It follows that BS-convergence implies $\mathrm{FO}_1^{\mathrm{local}}$-convergence.

Assume $(G_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}_1^{\mathrm{local}}$-convergent and let $\phi\in\mathrm{FO}_p^{\mathrm{local}}$ be an r-local formula. Let $\mathcal{F}_\phi$ be the set of all p-tuples $((F_1,f_1),\dots,(F_p,f_p))$ of rooted connected graphs with maximum degree at most d and radius (from the root) at most r such that $\bigcup_i F_i\models\phi(f_1,\dots,f_p)$.

Then, for every graph G the sets

$$\Omega_\phi(G)=\{(v_1,\dots,v_p):\ G\models\phi(v_1,\dots,v_p)\}$$

and

$$\biguplus_{((F_1,f_1),\dots,(F_p,f_p))\in\mathcal{F}_\phi}\prod_{i=1}^{p}\{v:\ G\models\theta_{(F_i,f_i)}(v)\}$$

differ by at most $O(|G|^{p-1})$ elements. Indeed, according to the definition of an r-local formula, the p-tuples $(x_1,\dots,x_p)$ belonging to exactly one of these sets are such that there exists $1\le i<j\le p$ such that $\mathrm{dist}(x_i,x_j)\le 2r$.

It follows that

$$\langle\phi,G\rangle=\Big(\sum_{((F_i,f_i))_{1\le i\le p}\in\mathcal{F}_\phi}\prod_{i=1}^{p}\langle\theta_{(F_i,f_i)},G\rangle\Big)+O(|G|^{-1}).$$

It follows that $\mathrm{FO}_1^{\mathrm{local}}$-convergence (hence BS-convergence) implies full $\mathrm{FO}^{\mathrm{local}}$-convergence. □

Remark 2.16. According to this proposition and Theorem 2.7, the BS-limit of a sequence of graphs with maximum degree at most D corresponds to a probability measure on $S(\mathcal{B}(\mathrm{FO}_1^{\mathrm{local}}))$ whose support is include in the clopen set $K(\zeta_D)$, where ζ_D is the sentence expressing that the maximum degree is at most D. The Boolean algebra $\mathcal{B}(\mathrm{FO}_1^{\mathrm{local}})$ is isomorphic to the Boolean algebra defined by the fragment $X\subset\mathrm{FO}_0(\lambda_1)$ of sentences for rooted graphs that are local with respect to the root (here, λ_1 denotes the signature of graphs augmented by one symbol of constant). According to this locality, any two countable rooted graphs (G_1,r_1) and (G_2,r_2), the trace of the complete theories of (G_1,r_1) and (G_2,r_2) on X are the same if and only if the (rooted) connected component (G_1',r_1) of (G_1,r_1) containing the root r_1 is elementary equivalent to the (rooted) connected component (G_2',r_2) of (G_2,r_2) containing the root r_2. As isomorphism and elementary equivalence are equivalent for countable connected graphs with bounded degrees (see Lemma 2.18) it is easily checked that $K_X(\zeta_D)$ is homeomorphic to $\mathcal{G}_D$. Hence our setting (while based on a very different and dual approach) leads essentially to the same limit object as [8] for BS-convergent sequences.

2.2.3. Elementary-convergence and FO_0-convergence. We already mentioned that FO_0-convergence is nothing but elementary convergence. Elementary convergence is implicitly part of the classical model theory. Although we only consider graphs here, the definition and results indeed generalize to general λ-structures We now reword the notion of elementary convergence:

A sequence $(G_n)_{n\in\mathbb{N}}$ is *elementarily convergent* if, for every sentence $\phi \in \mathrm{FO}_0$, there exists a integer N such that either all the graphs G_n $(n \geq N)$ satisfy ϕ or none of them do.

Of course, the limit object (as a graph) is not unique in general and formally, the limit of an elementarily convergent sequence of graphs is an elementary class defined by a complete theory.

Elementary convergence is also the backbone of all the X-convergences we consider in this paper. The FO_0-convergence is induced by an easy ultrametric defined on equivalence classes of elementarily equivalent graphs. Precisely, two (finite or infinite) graphs G_1, G_2 are *elementarily equivalent* (denoted $G_1 \equiv G_2$) if, for every sentence ϕ it holds

$$G_1 \models \phi \iff G_2 \models \phi.$$

In other words, two graphs are elementarily equivalent if they satisfy the same sentences.

A weaker (parametrized) notion of equivalence will be crucial: two graphs G_1, G_2 are *k-elementarily equivalent* (denoted $G_1 \equiv^k G_2$) if, for every sentence ϕ with quantifier rank at most k it holds $G_1 \models \phi \iff G_2 \models \phi$.

It is easily checked that for every two graphs G_1, G_2 it holds:

$$G_1 \equiv G_2 \qquad \iff \qquad (\forall k \in \mathbb{N})\ G_1 \equiv^k G_2.$$

For every fixed $k \in \mathbb{N}$, checking whether two graphs G_1 and G_2 are k-elementarily equivalent can be done using the so-called Ehrenfeucht-Fraïssé game.

¿From the notion of k-elementary equivalence naturally derives a pseudometric $\mathrm{dist}_0(G_1, G_2)$:

$$\mathrm{dist}_0(G_1, G_2) = \begin{cases} 0 & \text{if } G_1 \equiv G_2 \\ \min\{2^{-\mathrm{qrank}(\phi)} : (G_1 \models \phi) \wedge (G_2 \models \neg\phi)\} & \text{otherwise} \end{cases}$$

Proposition 2.2. *The metric space of countable graphs (up to elementary equivalence) with ultrametric* dist_0 *is compact.*

Proof. This is a direct consequence of the compactness theorem for first-order logic (a theory has a model if and only if every finite subset of it has a model) and of the downward Löwenheim-Skolem theorem (if a theory has a model and the language is countable then the theory has a countable model). □

Note that not every countable graph is (up to elementary equivalence) the limit of a sequence of finite graphs. A graph G that is a limit of a sequence finite graphs is said to have the *finite model property*, as such a graph is characterized by the property that every finite set of sentences satisfied by G has a finite model (what does not imply that G is elementarily equivalent to a finite graph).

Example 2.17. A *ray* is not an elementary limit of finite graphs as it contains exactly one vertex of degree 1 and all the other vertices have degree 2, what can be expressed in first-order logic but is satisfied by no finite graph. However, the union of two rays is an elementary limit from the sequence $(P_n)_{n\in\mathbb{N}}$ of paths of order n.

Although two finite graphs are elementary equivalent if and only if they are isomorphic, this property does not holds in general for countable graphs. For instance, the union of a ray and a line is elementarily equivalent to a ray. However we shall make use of the equivalence of isomorphisms and elementary equivalences for rooted connected countable locally finite graphs, which we prove now for completeness.

Lemma 2.18. *Let (G, r) and (G', r') be two rooted connected countable graphs. If G is locally finite then $(G, r) \equiv (G', r')$ if and only if (G, r) and (G', r') are isomorphic.*

Proof. If two rooted graphs are isomorphic they are obviously elementarily equivalent. Assume that (G, r) and (G', r') are elementarily equivalent. Enumerate the vertices of G in a way that distance to the root is not decreasing. Using n-back-and-forth equivalence (for all $n \in \mathbb{N}$), one builds a tree of partial isomorphisms of the subgraphs induced by the n first vertices, where ancestor relation is restriction. This tree is infinite and has only finite degrees. Hence, by Kőnig's lemma, it contains an infinite path. It is easily checked that it defines an isomorphism from (G, r) to (G', r') as these graphs are connected. $\square$

Fragments of FO_0 allow to define convergence notions, which are weaker that elementary convergence. The hierarchy of the convergence schemes defined by sub-algebras of $\mathcal{B}(\mathrm{FO}_0)$ is as strict as one could expect. Precisely, if $X \subset Y$ are two sub-algebras of $\mathcal{B}(\mathrm{FO}_0)$ then Y-convergence is strictly stronger than X-convergence — meaning that there exists graph sequences that are X-convergent but not Y-convergent — if and only if there exists a sentence $\phi \in Y$ such that for every sentence $\psi \in X$, there exists a (finite) graph G disproving $\phi \leftrightarrow \psi$.

We shall see that the special case of elementary convergent sequences is of particular importance. Indeed, every limit measure is a Dirac measure concentrated on a single point of $S(\mathcal{B}(\mathrm{FO}_0))$. This point is the complete theory of the elementary limit of the considered sequence. This limit can be represented by a finite or countable graph. As FO-convergence (and any FO_p-convergence) implies FO_0-convergence, the support of a limit measure μ corresponding to an FO_p-convergent sequence (or to an FO-convergent sequence) is such that $\mathrm{Supp}(\mu)$ projects to a single point of $S(\mathcal{B}(\mathrm{FO}_0))$.

Finally, let us remark that all the results of this section can be readily formulated and proved for λ-structures.

2.3. Combining Fragments

2.3.1. The FO_p Hierarchy. When we consider FO_p-convergence of finite λ-structures for finite a signature λ, the space $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ can be given the following ultrametric dist_p (compatible with the topology of $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$): Let $T_1, T_2 \in S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ (where the points of $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ are identified with ultrafilters on $\mathcal{B}(\mathrm{FO}_p(\lambda))$). Then

$$\mathrm{dist}_p(T_1, T_2) = \begin{cases} 0 & \text{if } T_1 = T_2 \\ 2^{-\min\{\mathrm{qrank}(\phi):\ \phi \in T_1 \setminus T_2\}} & \text{otherwise} \end{cases}$$

This ultrametric has several other nice properties:

- actions of S_p on $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ are isometries:

$$\forall \sigma \in S_p\ \forall T_1, T_2 \in S(\mathcal{B}(\mathrm{FO}_p(\lambda))) \quad \mathrm{dist}_p(\sigma \cdot T_1, \sigma \cdot T_2) = \mathrm{dist}_p(T_1, T_2);$$

- projections π_p are contractions:

$$\forall q \geq p\ \forall T_1, T_2 \in S(\mathcal{B}(\mathrm{FO}_q(\lambda))) \quad \mathrm{dist}_p(\pi_p(T_1), \pi_p(T_2)) \leq \mathrm{dist}_q(T_1, T_2);$$

We prove that there is a natural isometric embedding $\eta_p : S(\mathcal{B}(\mathrm{FO}_p(\lambda))) \to S(\mathcal{B}(\mathrm{FO}(\lambda)))$. This may be seen as follows: for an ultrafilter $X \in S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$, consider the filter X^+ on $\mathcal{B}(\mathrm{FO}(\lambda))$ generated by X and all the formulas $x_i = x_{i+1}$ (for $i \geq p$). This filter is an ultrafilter: for every sentence $\phi \in \mathrm{FO}(\lambda)$, let $\widetilde{\phi}$ be the sentence obtained from ϕ by replacing each free occurrence of a variable x_q with $q > p$ by x_p. It is clear that ϕ and $\widetilde{\phi}$ are equivalent modulo the theory $T_p = \{(x_i = x_{i+1}) :\ i \geq p\}$. As either $\widetilde{\phi}$ or $\neg\widetilde{\phi}$ belongs to X, either ϕ or $\neg\phi$ belongs to $\eta_p(X)$. Moreover, we deduce easily from the fact that $\widetilde{\phi}$ and ϕ have the same quantifier rank that if $q \geq p$ then $\pi_q \circ \eta_p$ is an isometry. Finally, let us note that $\pi_p \circ \eta_p$ is the identity of $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$.

Let λ_p be the signature λ augmented by p symbols of constants $c_1, \dots, c_p$. There is a natural isomorphism of Boolean algebras $\nu_p : \mathrm{FO}_p(\lambda) \to \mathrm{FO}_0(\lambda_p)$, which replaces the free occurrences of the variables $x_1, \dots, x_p$ in a formula $\phi \in \mathrm{FO}_p$ by the corresponding symbols of constants $c_1, \dots, c_p$, so that it holds, for every modeling $\mathbf{A}$, for every $\phi \in \mathrm{FO}_p$ and every $v_1, \dots, v_p \in A$:

$$\mathbf{A} \models \phi(v_1, \dots, v_p) \qquad \iff \qquad (\mathbf{A}, v_1, \dots, v_p) \models \nu_p(\phi).$$

This mapping induces an isometric isomorphism of the metric spaces $(S(\mathcal{B}(\mathrm{FO}_p(\lambda))), \mathrm{dist}_p)$ and $(S(\mathcal{B}(\mathrm{FO}_0(\lambda_p))), \mathrm{dist}_0)$. Note that the Stone space $S(\mathcal{B}(\mathrm{FO}_0(\lambda_p)))$ associated to the Boolean algebra $\mathcal{B}(\mathrm{FO}_0(\lambda_p))$ is the space of all complete theories of λ_p-structures. In particular, points of $S(\mathcal{B}(\mathrm{FO}_p(\lambda))$ can be represented (up to elementary equivalence) by countable λ-structures with p special points. All these transformations may seem routine but they need to be carefully formulated and checked.

We can test whether the distance dist_p of two theories T and T' is smaller than 2^{-n} by means of an Ehrenfeucht-Fraïssé game: Let $\nu_p(T) = \{\nu_p(\phi) : \phi \in T\}$ and, similarly, let $\nu_p(T') = \{\nu_p(\phi) : \phi \in T'\}$. Let $(\mathbf{A}, v_1, \dots, v_p)$ be a model of T and let $(\mathbf{A}', v_1', \dots, v_p')$ ba a model of T'. Then it holds

$$\mathrm{dist}_p(T, T') < 2^{-n} \qquad \iff \qquad (\mathbf{A}, v_1, \dots, v_p) \equiv^n (\mathbf{A}', v_1', \dots, v_p').$$

Recall that the n-rounds *Ehrenfeucht-Fraïssé game* on two λ-structures $\mathbf{A}$ and $\mathbf{A}'$, denoted $\mathrm{EF}(\mathbf{A}, \mathbf{A}', n)$ is the perfect information game with two players — the Spoiler and the Duplicator — defined as follows: The game has n rounds and each round has two parts. At each round, the Spoiler first chooses one of $\mathbf{A}$ and $\mathbf{A}'$ and accordingly selects either a vertex $x \in A$ or a vertex $y \in A'$. Then, the Duplicator selects a vertex in the other λ-structure. At the end of the n rounds, n vertices have been selected from each structure: $x_1, \dots, x_n$ in A and $y_1, \dots, y_n$ in A' (x_i and y_i corresponding to vertices x and y selected during the ith round). The Duplicator wins if the substructure induced by the selected vertices are order-isomorphic (i.e. $x_i \mapsto y_i$ is an isomorphism of $\mathbf{A}[\{x_1, \dots, x_n\}]$ and $\mathbf{A}'[\{y_1, \dots, y_n\}]$). As there are no hidden moves and no draws, one of the two players has a winning strategy, and we say that that player wins $\mathrm{EF}(\mathbf{A}, \mathbf{A}', n)$. The main property of this game is the following equivalence, due to Fraïssé [32, 33] and Ehrenfeucht [27]: The duplicator wins $\mathrm{EF}(\mathbf{A}, \mathbf{A}', n)$ if and only if $\mathbf{A} \equiv^n \mathbf{A}'$. In our context this translates to the following equivalence:

$$\mathrm{dist}_p(T, T') < 2^{-n} \iff \text{Duplicator wins } \mathrm{EF}((\mathbf{A}, v_1, \dots, v_p), (\mathbf{A}', v_1', \dots, v_p'), n).$$

As $\mathrm{FO}_0 \subset \mathrm{FO}_1 \subset \cdots \subset \mathrm{FO}_p \subset \mathrm{FO}_{p+1} \subset \cdots \subset \mathrm{FO} = \bigcup_i \mathrm{FO}_i$, the fragments FO form a hierarchy of more and more restrictive notions of convergence. In particular, FO_{p+1}-convergence implies FO_p-convergence and FO-convergence is equivalent to FO_p for all p. If a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is FO_p-convergent then for every $q \leq p$ the FO_q-limit of $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is a measure $\mu_q \in \mathrm{rca}(S(\mathcal{B}(\mathrm{FO}_q)))$, which is the pushforward of μ_p by the projection π_q (more precisely, by the restriction of π_q to $S(\mathcal{B}(\mathrm{FO}_p)))$:

$$\mu_q = (\pi_q)_*(\mu_p).$$

2.3.2. $\mathrm{FO}^{\mathrm{local}}$ **and Locality.** FO-convergence can be reduced to the conjunction of elementary convergence and $\mathrm{FO}^{\mathrm{local}}$-convergence, which we call *local convergence*. This is a consequence of a result, which we recall now:

THEOREM 2.19 (Gaifman locality theorem [35]). *For every first-order formula $\phi(x_1, \dots, x_n)$ there exist integers t and r such that ϕ is equivalent to a Boolean combination of t-local formulas $\xi_s(x_{i_1}, \dots, x_{i_s})$ and sentences of the form*

$$\exists y_1 \dots \exists y_m \Big(\bigwedge_{1\leq i<j\leq m} \mathrm{dist}(y_i, y_j) > 2r \wedge \bigwedge_{1\leq i\leq m} \psi(y_i) \Big) \tag{2.1}$$

where ψ is r-local. Furthermore, we can choose

$$r \leq 7^{\mathrm{qrank}(\phi)-1},\ t \leq (7^{\mathrm{qrank}(\phi)-1} - 1)/2,\ m \leq n + \mathrm{qrank}(\phi),$$

and, if ϕ is a sentence, only sentences (2.1) occur in the Boolean combination. Moreover, these sentences can be chosen with quantifier rank at most $q(\mathrm{qrank}(\phi))$, for some fixed function q.

¿From this theorem and the following folklore technical result will follow the claimed decomposition of FO-convergence into elementary and local convergence.

LEMMA 2.20. *Let B be a Boolean algebra, let A_1 and A_2 be sub-Boolean algebras of B, and let $b \in B[A_1 \cup A_2]$ be a Boolean combination of elements from A_1 and A_2. Then b can be written as*

$$b = \bigvee_{i\in I} x_i \wedge y_i,$$

where I is finite, $x_i \in A_1$, $y_i \in A_2$, and for every $i \neq j$ in I it holds $(x_i \wedge y_i) \wedge (x_j \wedge y_j) = 0$.

PROOF. Let $b = F(u_1, \dots, u_a, v_1, \dots, v_b)$ with $u_i \in A_1$ $(1 \leq i \leq a)$ and $v_j \in A_2$ $(1 \leq j \leq b)$ where F is a Boolean combination. By using iteratively Shannon's expansion, we can write F as

$$F(u_1, \dots, u_a, v_1, \dots, v_b) = \bigvee_{(X_1,X_2,Y_1,Y_2)\in\mathcal{F}} \Big(\bigwedge_{i\in X_1} u_i \wedge \bigwedge_{i\in X_2} \neg u_i \wedge \bigwedge_{j\in Y_1} v_j \wedge \bigwedge_{j\in Y_2} \neg v_j\Big),$$

where $\mathcal{F}$ is a subset of the quadruples (X_1, X_2, Y_1, Y_2) such that (X_1, X_2) is a partition of $[a]$ and (Y_1, Y_2) is a partition of $[b]$. For a quadruple $Q = (X_1, X_2, Y_1, Y_2)$, define $x_Q = \bigwedge_{i\in X_1} u_i \wedge \bigwedge_{i\in X_2} \neg u_i$ and $y_Q = \bigwedge_{j\in Y_1} v_j \wedge \bigwedge_{j\in Y_2} \neg v_j$. Then for every $Q \in \mathcal{F}$ it holds $x_Q \in A_1, y_Q \in A_2$, for every $Q \neq Q' \in \mathcal{F}$ it holds $x_Q \wedge y_Q \wedge x_{Q'} \wedge y_{Q'} = 0$, and we have $b = \bigvee_{Q\in\mathcal{F}} x_Q \wedge y_Q$. □

THEOREM 2.21. *Let $(\mathbf{A}_n)$ be a sequence of finite λ-structures. Then $(\mathbf{A}_n)$ is* FO*-convergent if and only if it is both* $\mathrm{FO}^{\mathrm{local}}$*-convergent and* FO_0*-convergent. Precisely, $(\mathbf{A}_n)$ is* FO_p*-convergent if and only if it is both* $\mathrm{FO}_p^{\mathrm{local}}$*-convergent and* FO_0*-convergent.*

PROOF. Assume $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is both $\mathrm{FO}_p^{\mathrm{local}}$-convergent and FO_0-convergent and let $\phi \in \mathrm{FO}_p$. According to Theorem 2.19, there exist integers t and r such that ϕ is equivalent to a Boolean combination of t-local formula $\xi(x_{i_1}, \dots, x_{i_s})$ and of sentences. As both $\mathrm{FO}^{\mathrm{local}}$ and FO_0 define a sub-Boolean algebra of $\mathcal{B}(\mathrm{FO})$, according to Lemma 2.20, ϕ can be written as $\bigvee_{i\in I} \psi_i \wedge \theta_i$, where I is finite, $\psi_i \in \mathrm{FO}^{\mathrm{local}}$, $\theta_i \in \mathrm{FO}_0$, and $\psi_i \wedge \theta_i \wedge \psi_j \wedge \theta_j = 0$ if $i \neq j$. Thus for every finite λ-structure $\mathbf{A}$ it holds

$$\langle \phi, \mathbf{A}\rangle = \sum_{i\in I} \langle \psi_i \wedge \theta_i, \mathbf{A}\rangle.$$

As $\langle \cdot, \mathbf{A}\rangle$ is additive and $\langle \theta_i, \mathbf{A}\rangle \in \{0, 1\}$ we have $\langle \psi_i \wedge \theta_i, \mathbf{A}\rangle = \langle \psi_i, \mathbf{A}\rangle \, \langle \theta_i, \mathbf{A}\rangle$. Hence

$$\langle \phi, \mathbf{A}\rangle = \sum_{i\in I} \langle \psi_i, \mathbf{A}\rangle \, \langle \theta_i, \mathbf{A}\rangle.$$

Thus if $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is both $\mathrm{FO}_p^{\mathrm{local}}$-convergent and FO_0-convergent then $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is FO_p-convergent. □

Similarly that points of $S(\mathcal{B}(\mathrm{FO}_p(\lambda))$ can be represented (up to elementary equivalence) by countable λ-structures with p special points, points of $S(\mathcal{B}(\mathrm{FO}_p^{\mathrm{local}}(\lambda))$ can be represented by countable λ-structures with p special points such that every connected component contains at least one special point. In particular, points of $S(\mathcal{B}(\mathrm{FO}_1^{\mathrm{local}}(\lambda))$ can be represented by rooted connected countable λ-structures.

Also, the structure of an $\mathrm{FO}_2^{\mathrm{local}}$-limit of graphs can be outlined by considering that points of $S(\mathcal{B}(\mathrm{FO}_2^{\mathrm{local}}))$ as countable graphs with two special vertices c_1 and c_2, such that every connected component contains at least one of c_1 and c_2. Let μ_2 be the limit probability measure on $S(\mathcal{B}(\mathrm{FO}_2^{\mathrm{local}}))$ for an $\mathrm{FO}_2^{\mathrm{local}}$-convergent sequence $(G_n)_{n\in\mathbb{N}}$, let π_1 be the standard projection of $S(\mathcal{B}(\mathrm{FO}_2^{\mathrm{local}}))$ into $S(\mathcal{B}(\mathrm{FO}_1^{\mathrm{local}}))$, and let μ_1 be the pushforward of μ_2 by π_1. We construct a measurable graph $\hat{G}$ as follows: the vertex set of $\hat{G}$ is the support $\mathrm{Supp}(\mu_1)$ of μ_1. Two vertices x and y of $\hat{G}$ are adjacent if there exists $x' \in \pi_1^{-1}(x)$ and $y' \in \pi_1^{-1}(y)$ such that (considered as ultrafilters of $\mathcal{B}(\mathrm{FO}_2^{\mathrm{local}})$) it holds:

- $x_1 \sim x_2$ belongs to both x' and y',
- the transposition $\tau_{1,2}$ exchanges x' and y' (i.e. $y' = \tau_{1,2} \cdot x'$).

The vertex set of $\hat{G}$ is of course endowed with a structure of a probability space (as a measurable subspace of $S(\mathcal{B}(\mathrm{FO}_1^{\mathrm{local}}))$ equipped with the probability measure μ_1). In the case of bounded degree graphs, the obtained graph $\hat{G}$ is the *graph of graphs* introduced in [52]. Notice that this graph may have loops. An example of such a graph is shown Fig. 1.

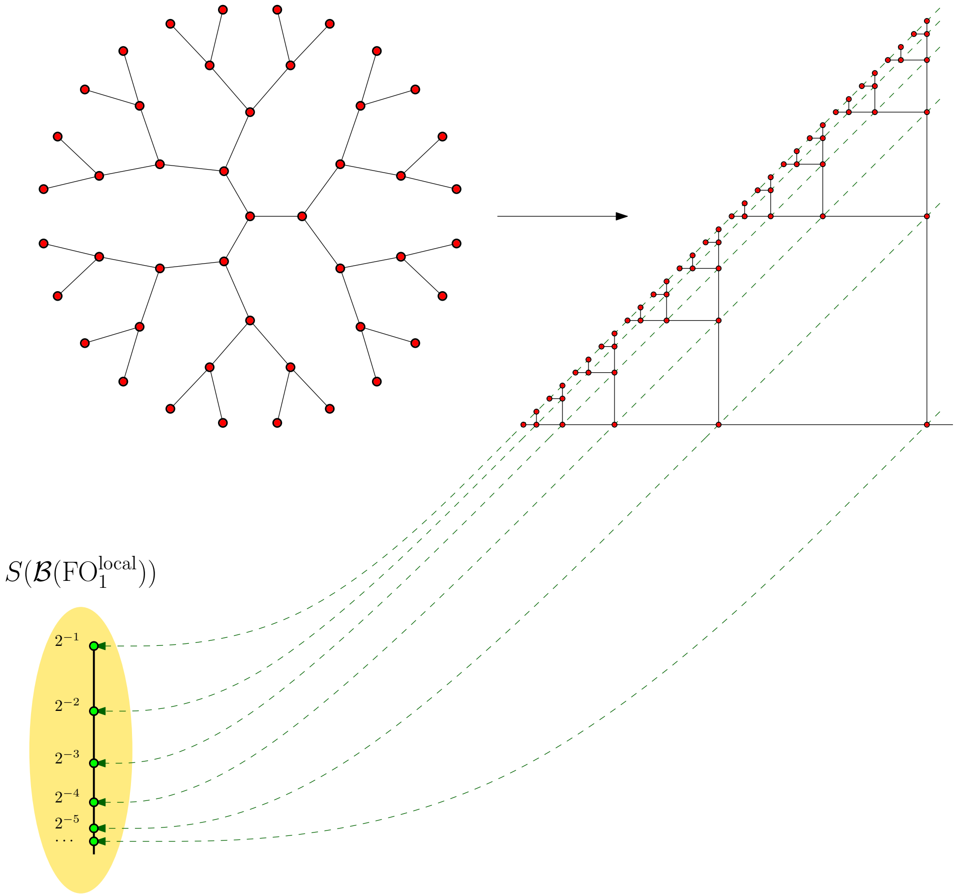

FIGURE 1. An outline of the local limit of a sequence of trees

2.3.3. Component-Local Formulas. It is sometimes possible to reduce $\mathrm{FO}_{\mathrm{local}}$ to a smaller fragment. This is in particular the case when connected components of the considered structures can be identified by some first-order formula. Precisely:

DEFINITION 2.22. Let λ be a signature and let T be a theory of λ-structures. A binary relation $\varpi \in \lambda$ is a *component relation* in T if T entails that ϖ is an equivalence relation such that for every k-ary relation $R \in \lambda$ with $k \geq 2$ it holds

$$T \models (\forall x_1, \dots, x_k) \Big(R(x_1, \dots, x_k) \to \bigwedge_{1 \leq i < j \leq k} \varpi(x_i, x_j) \Big).$$

A local formula ϕ with p free variables is ϖ*-local* if ϕ is equivalent (modulo T) to $\phi \wedge \bigwedge_{x_i, x_j \in \mathrm{Fv}(\phi)} \varpi(x_i, x_j)$.

In presence of a component relation, it is possible to reduce from $\mathrm{FO}_{\mathrm{local}}$ to the fragment of ϖ-local formulas, thanks to the following result.

LEMMA 2.23. *Let ϖ be a component relation in a theory T. For every local formula ϕ with quantifier rank r there exist ϖ-local formulas $\xi_{i,j} \in \mathrm{FO}^{\mathrm{local}}_{q_{i,j}}$ $(1 \leq i \leq n,\ j \in I_i)$ with quantifier rank at most r and permutations σ_i of $[p]$ $(1 \leq i \leq n)$ such that for each $1 \leq i \leq n$, $\sum_{j \in I_i} q_{i,j} = p$ and, for every model $\mathbf{A}$ of T it holds*

$$\Omega_\phi(\mathbf{A}) = \biguplus_{i=1}^{n} F_{\sigma_i}\Big(\prod_{j \in I_i} \Omega_{\xi_{i,j}}(\mathbf{A}) \Big),$$

where $F_{\sigma_i}(X)$ performs a permutation of the coordinates according to σ_i.

Proof. First note that if two ϖ-local formulas ϕ_1 and ϕ_2 share a free variable then $\phi_1 \wedge \phi_2$ is ϖ-local. For this obvious fact, we deduce that if $\psi_1, \dots, \psi_n$ are ϖ-local formulas in FO_p, then there is a partition τ and a permtutation σ of $[p]$ such that for every λ-structure $\mathbf{A}$ it holds

$$\Omega_{\bigwedge_{i=1}^n \psi_i}(\mathbf{A}) = F_\sigma\Big(\prod_{P\in\tau} \Omega_{\bigwedge_{i\in P}\psi_i}(\mathbf{A})\Big),$$

where each $\bigwedge_{i\in P}\psi_i$ is ϖ-local, and $F_\sigma : A^p \to A^p$ is defined by

$$F_\sigma(X) = \{(v_{\sigma(1)}, \dots, v_{\sigma(p)}) :\ (v_1, \dots, v_p) \in X\}.$$

For a partition τ of $[p]$ we denote by ζ_τ the conjunction of $\varpi(x_i, x_j)$ for every i, j belonging to a same part and of $\neg\varpi(x_i, x_j)$ for every i, j belonging to different parts. Then, for any two distinct partitions τ and τ', the formula $\zeta_\tau \wedge \zeta'_\tau$ is never satisfied; moreover $\bigvee_\tau \zeta_\tau$ is always satisfied. Thus for every local formula ϕ it holds

$$\phi = \bigvee_\tau (\zeta_\tau \wedge \phi) = \bigoplus_\tau (\zeta_\tau \wedge \phi)$$

(where only the partitions τ for which $\zeta_\tau \wedge \phi \neq 0$ have to be considered).

We denote by Λ_τ the formula $\bigwedge_{P\in\tau}\bigwedge_{i,j\in P}\varpi(x_i, x_j)$. Obviously it holds

$$\Lambda_\tau = \bigoplus_{\tau'\geq\tau} \zeta_\tau,$$

where $\oplus$ stands for the exclusive disjunction ($a \oplus b = (a \wedge \neg n) \vee (\neg a \wedge b)$) and $\tau' \geq \tau$ means that τ' is a partition of $[p]$, which is coarser than τ. Then there exists (by Möbius inversion or immediate induction) a function M from the set of the partitions of $[p]$ to the powerset of the set of paritions of $[p]$ such that for every partition τ of $[p]$ it holds

$$\zeta_\tau = \bigoplus_{\tau'\in M(\tau)} \Lambda_{\tau'}.$$

Hence

$$\phi = \bigoplus_\tau \bigoplus_{\tau'\in M(\tau)} \Lambda_{\tau'} \wedge \phi.$$

It follows that ϕ is a Boolean combination of formulas $\Lambda_\tau \wedge \phi$, for partitions τ such that $\zeta_\tau \wedge \phi \neq 0$ (as $\zeta_\tau \wedge \phi \neq 0$ and $\tau' \geq \tau$ imply $\zeta_{\tau'} \wedge \phi \neq 0$). Each formula $\Lambda_\tau \wedge \phi$ is itself a Boolean combination of ϖ-local formulas. Putting this in standard form (exclusive disjunction of conjunctions) and gathering in the conjunctions the ϖ-local formulas whose set of free variables intersect, we get that there exists families $\mathcal{F}_\tau$ of ϖ-local formulas φ_P ($P \in \tau$) with free variables $\mathrm{Fv}(\varphi_P) = \{x_j : j \in P\}$ such that

$$\phi = \bigvee_\tau \bigvee_{\varphi\in\mathcal{F}_\tau} \bigwedge_{P\in\tau} \varphi_P,$$

where the disjunction is exclusive.

Hence, considering adequate permutations σ_τ of $[p]$ it holds

$$\Omega_\phi(\mathbf{A}) = \biguplus_\tau \biguplus_{\varphi\in\mathcal{F}_\tau} F_{\sigma_\tau}\Big(\prod_{P\in\tau} \Omega_{\tilde{\varphi}_P}(\mathbf{A})\Big),$$

which is the requested form.

Note that the fact that $\mathrm{qrank}(\xi_{i,j}) \leq \mathrm{qrank}(\phi)$ is obvious as we did not introduce qny quantifier in our transformations. □

As a consequence, we get the the desired:

Corollary 2.1. *Let ϖ be a component relation in a theory T and let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be a sequence of models of T. Then the sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}_{\mathrm{local}}$-convergent if and only if it is $\mathrm{FO}_{\varpi\text{-}local}$-convergent.*

2.3.4. Sequences with Homogeneous Elementary Limit. Elementary convergence is an important aspect of FO-convergence and we shall see that in several contexts, FO-convergence can be reduced to the conjunction and elementary convergence of X-convergence (for some suitable fragment X).

In some special cases, the limit (as a countable structure) will be unique. This means that some particular complete theories have exactly one countable model (up to isomorphism). Such complete theories are called *ω-categorical.* Several properties are known to be equivalent to ω-categoricity. For instance, for a complete theory T the following statements are equivalent:

- T is ω-categorical;
- for every every $p \in \mathbb{N}$, the Stone space $S(\mathcal{B}(\mathrm{FO}_p(\lambda), T))$ is finite (see Fig. 2);
- every countable model $\mathbf{A}$ of T has an *oligomorphic* automorphism group, what means that for every $n \in \mathbb{N}$, A^n has finitely many orbits under the action of $\mathrm{Aut}(\mathbf{A})$.

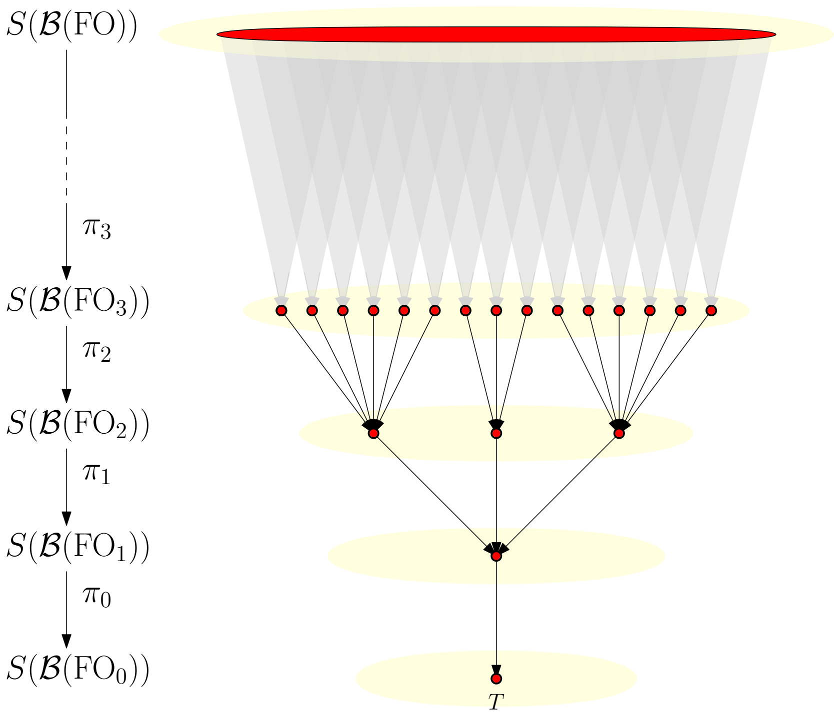

Figure 2. Ultrafilters projecting to an ω-categorical theory

A theory T is said to have *quantifier elimination* if, for every formula $\phi \in \mathrm{FO}_p(\lambda)$ there exists $\widetilde{\phi} \in \mathrm{QF}_p(\lambda)$ such that $T \models \phi \leftrightarrow \widetilde{\phi}$. If a theory has quantifier elimination then it is ω-categorical. Indeed, for every p, there exists only finitely many quantifier free formulas with p free variables hence (up to equivalence modulo T) only finitely many formulas with p free variables. The unique countable model of a complete theory T with quantifier elimination is *ultra-homogeneous*, what means that every partial isomorphism of finite induced substructures extends as a full automorphism. In the context of relational structures, the property of having a countable ultra-homogeneous model is equivalent to the property of having quantifier elimination. We provide a proof of this folklore result (in the context of graphs) in order to illustrate these notions.

LEMMA 2.24. *Let T be a complete theory (of graphs) with no finite model.*
Then T has quantifier elimination if and only if some (equivalently, every) countable model of T is ultra-homogeneous.

PROOF. Assume that T has an ultra-homogeneous countable model G. Let $(a_1, \ldots, a_p)$, $(b_1, \ldots, b_p)$ be p-tuples of vertices of G. Assume that $a_i \mapsto b_i$ is an isomorphism between $G[a_1, \ldots, a_p]$ and $G[b_1, \ldots, b_p]$. Then, as G is ultra-homogeneous, there exists an automorphism f of G such that $f(a_i) = b_i$ for every $1 \leq i \leq p$. As the satisfaction of a first-order formula is invariant by the action of the automorphism group, for every formula $\phi \in \mathrm{FO}_p$ it holds

$$G \models \phi(a_1, \ldots, a_p) \quad \iff \quad G \models \phi(b_1, \ldots, b_p).$$

Consider a maximal set $\mathcal{F}$ of p-tuples $(v_1, \ldots, v_p)$ of G such that $G \models \phi(v_1, \ldots, v_p)$ and no two p-tuples induce isomorphic (ordered) induced subgraphs. Obviously $|\mathcal{F}| = 2^{O(p^2)}$ is finite. Moreover, each p-tuple $\vec{v} = (v_1, \ldots, v_p)$ defines a quantifier free formula $\eta_{\vec{v}}$ with p free variables such that $G \models \eta_{\vec{v}}(x_1, \ldots, x_p)$ if and only if $x_i \mapsto v_i$ is an isomorphism between $G[x_1, \ldots, x_p]$ and $G[v_1, \ldots, v_p]$. Hence it holds:

$$G \models \phi \leftrightarrow \bigvee_{\vec{v} \in \mathcal{F}} \eta_{\vec{v}}.$$

In other words, ϕ is equivalent (modulo T) to the quantifier free formula $\widetilde{\phi} = \bigvee_{\vec{v} \in \mathcal{F}} \eta_{\vec{v}}$, that is: T has quantifier elimination.

Conversely, assume that T has quantifier elimination. As notice above, T is ω-categorical thus has a unique countable model. Assume $(a_1, \ldots, a_p)$ and $(b_1, \ldots, b_p)$ are p-tuples of vertices such that $f : a_i \mapsto b_i$ is a partial isomorphism. Assume that f does not extend into an automorphism of G. Let $(a_1, \ldots, a_q)$ be a tuple of vertices of G of maximal length such that there exists $b_{p+1}, \ldots, b_q$ such that $a_i \mapsto b_i$ is a partial isomorphism. Let a_{q+1} be a vertex distinct from $a_1, \ldots, a_q$. Let $\phi(x_1, \ldots, x_q)$ be the formula

$$\begin{aligned} &\bigwedge_{a_i \sim a_j} (x_i \sim x_j) \wedge \bigwedge_{\neg(a_i \sim a_j)} \neg(x_i \sim x_j) \wedge \bigwedge_{1 \leq i \leq q} \neg(x_i = x_j) \\ \wedge\ (\exists y) &\left(\bigwedge_{a_i \sim a_{q+1}} (x_i \sim y) \wedge \bigwedge_{\neg(a_i \sim a_{q+1})} \neg(x_i \sim y) \wedge \bigwedge_{1 \leq i \leq q} \neg(x_i = y) \right) \end{aligned}$$

As T has quantifier elimination, there exists a quantifier free formula $\widetilde{\phi}$ such that $T \models \phi \leftrightarrow \widetilde{\phi}$. As $G \models \phi(a_1, \ldots, a_q)$ (witnessed by a_{q+1}) it holds $G \models \widetilde{\phi}(a_1, \ldots, a_q)$ hence $G \models \widetilde{\phi}(b_1, \ldots, b_q)$ (as $a_i \mapsto b_i, 1 \leq i \leq q$ is a partial isomorphism) thus $G \models \phi(b_1, \ldots, b_q)$. It follows that there exists b_{q+1} such that $a_i \mapsto b_i, 1 \leq i \leq q+1$ is a partial isomorphism, contradicting the maximality of $(a_1, \ldots, a_q)$. $\square$

When a sequence of graphs is elementarily convergent to an ultra-homogeneous graph (i.e. to a complete theory with quantifier elimination), we shall prove that FO-convergence reduces to QF-convergence. This later mode of convergence is of particular interest as it is equivalent to L-convergence.

LEMMA 2.25. *Let $(G_n)_{n\in\mathbb{N}}$ be sequence of graphs that converges elementarily to some ultra-homogeneous graph $\hat{G}$. Then the following properties are equivalent:*
- *the sequence $(G_n)_{n\in\mathbb{N}}$ is* FO*-convergent;*
- *the sequence $(G_n)_{n\in\mathbb{N}}$ is* QF*-convergent.*
- *the sequence $(G_n)_{n\in\mathbb{N}}$ is L-convergent.*

PROOF. As FO-convergence implies QF-convergence we only have to prove the opposite direction. Assume that the sequence $(G_n)_{n\in\mathbb{N}}$ is QF-convergent. According to Lemma 2.24, for every formula $\phi \in \mathrm{FO}_p$ there exists a quantifier free formula $\widetilde{\phi} \in \mathrm{QF}_p$ such that $\hat{G} \models \phi \leftrightarrow \widetilde{\phi}$ (i.e. $\mathrm{Th}(\hat{G})$ has quantifier elimination). As $\hat{G}$ is an elementary limit of the sequence $(G_n)_{n\in\mathbb{N}}$ there exists N such that for every $n \geq N$ it holds $G_n \models \phi \leftrightarrow \widetilde{\phi}$. It follows that for every $n \geq N$ it holds $\langle \phi, G_n\rangle = \langle \widetilde{\phi}, G_n\rangle$ hence $\lim_{n\to\infty}\langle \phi, G_n\rangle$ exists. Thus the sequence $(G_n)_{n\in\mathbb{N}}$ is FO-convergent. □

There are not so many countable ultra-homogeneous graphs.

THEOREM 2.26 (Lachlan and Woodrow [47]). *Every infinite countable ultrahomogeneous undirected graph is isomorphic to one of the following:*
- *the disjoint union of m complete graphs of size n, where $m, n \leq \omega$ and at least one of m or n is ω, (or the complement of it);*
- *the generic graph for the class of all countable graphs not containing K_n for a given $n \geq 3$ (or the complement of it).*
- *the Rado graph R (the generic graph for the class of all countable graphs).*

Among them, the Rado graph R (also called "the random graph") is characterized by the *extension property*: for every finite disjoint subsets of vertices A and B of R there exists a vertex z of $R - A - B$ such that z is adjacent to every vertex in A and to no vertex in B. We deduce for instance the following application of Lemma 2.25.

EXAMPLE 2.27. It is known [9, 10] that for every fixed k, Paley graphs of sufficiently large order satisfy the k-extension property hence the sequence of Paley graphs converge elementarily to the Rado graph. Moreover, Paley graphs is a standard example of quasi-random graphs [19], and the sequence of Paley graphs is L-convergent to the 1/2-graphon. Thus, according to Lemma 2.25, the sequence of Paley graphs is FO-convergent.

We now relate more precisely the extension property with quantifier elimination.

DEFINITION 2.28. Let $k \in \mathbb{N}$. A graph G has the *k-extension property* if, for every disjoint subsets of vertices A, B of G with size k there exists a vertex z not in $A \cup B$ that is adjacent to every vertex in A and to no vertex in B. In other words,

G has the k-extension property if G satisfies the sentence Υ_k below:

$$(\forall x_1,\dots,x_{2k})\quad \Biggl(\bigwedge_{1\le i<j\le 2k}\neg(x_i=x_j)$$
$$\to\quad (\exists z)\ \bigwedge_{i=1}^{2k}\neg(x_i=z)\wedge\bigwedge_{i=1}^{k}(x_i\sim z)\wedge\bigwedge_{i=k+1}^{2k}\neg(x_i\sim z)\Biggr)$$

LEMMA 2.29. *Let G be a graph and let p,r be integers.*

If G has the $(p+r)$-extension property then every formula ϕ with p free variables and quantifier rank r is equivalent, in G, with a quantifier free formula.

PROOF. Let ϕ be a formula with p free variables and quantifier rank r. Let $(a_1,\dots,a_p)$ and $(b_1,\dots,b_p)$ be two p-tuples of vertices of G such that $a_i\mapsto b_i$ is a partial isomorphism. The $(p+r)$-extension properties allows to easily play a r-turns back-and-forth game between $(G,a_1,\dots,a_p)$ and $(G,b_1,\dots,b_p)$, thus proving that $(G,a_1,\dots,a_p)$ and $(G,b_1,\dots,b_p)$ are r-equivalent. It follows that $G\models\phi(a_1,\dots,a_p)$ if and only if $G\models\phi(b_1,\dots,b_p)$. Following the lines of Lemma 2.24, we deduce that there exists a quantifier free formula $\widetilde{\phi}$ such that $G\models\phi\leftrightarrow\widetilde{\phi}$. □

We now prove that random graphs converge elmentarily to the countable random graphs.

LEMMA 2.30. *Let $1/2>\delta>0$. Assume that for every positive integer $n\ge 2$ and every $1\le i<j\le n$, $p_{n,i,j}\in[\delta,1-\delta]$. Assume that for each $n\in\mathbb{N}$, G_n is a random graph on $[f(n)]$ where $f(n)\ge n$, and where i and j are adjacent with probability $p_{n,i,j}$ (all these events being independent). Then the sequence $(G_n)_{n\in\mathbb{N}}$ almost surely converges elementarily to the Rado graph.*

PROOF. Let $p\in\mathbb{N}$ and let $\alpha=\delta(1-\delta)$. The probability that $G_n\models\Upsilon_p$ is at least $1-(1-\alpha^p)^{f(n)}$. It follows that for $N\in\mathbb{N}$ the probability that all the graphs G_n $(n\ge N)$ satisfy Υ_p is at least $1-\alpha^{-p}(1-\alpha^p)^{f(N)}$. According to Borel-Cantelli lemma, the probability that G_n does not satisfy Υ_p infinitely many is zero. As this holds for every integer p, it follows that, with high probability, every elementarily converging subsequence of $(G_n)_{n\in\mathbb{N}}$ converges to the Rado graph hence, with high probability, $(G_n)_{n\in\mathbb{N}}$ converges elementarily to the Rado graph. □

Thus we get:

THEOREM 2.31. *Let $0<p<1$ and let $G_n\in G(n,p)$ be independent random graphs with edge probability p. Then $(G_n)_{n\in\mathbb{N}}$ is almost surely* FO*-convergent.*

PROOF. This is an immediate consequence of Lemma 2.25, Lemma 2.30 and the easy fact that $(G_n)_{n\in\mathbb{N}}$ is almost surely QF-convergent. □

THEOREM 2.32. *For every $\phi\in\mathrm{FO}_p$ there exists a polynomial $P_\phi\in\mathbb{Z}[X_1,\dots,X_{\binom{p}{2}}]$ such that for every sequence $(G_n)_{n\in\mathbb{N}}$ of finite graphs that converges elementarily to the Rado graph the following holds:*

If $(G_n)_{n\in\mathbb{N}}$ is L-convergent to some graphon W then

$$\lim_{n\to\infty}\langle\phi,G_n\rangle=\int\cdots\int P_\phi((W_{i,j}(x_i,x_j))_{1\le i<j\le p})\,\mathrm{d}x_1\dots\mathrm{d}x_p.$$

PROOF. Assume the sequence $(G_n)_{n\in\mathbb{N}}$ is elementarily convergent to the Rado graph and that it is L-convergent to some graphon W.

According to Lemma 2.24, there exists a quantifier free formula $\widetilde{\phi}$ such that

$$G \models (\forall x_1 \dots x_p)\ \phi(x_1, \dots, x_p) \leftrightarrow \widetilde{\phi}(x_1, \dots, x_p)$$

(hence $\Omega_\phi(G) = \Omega_{\widetilde{\phi}}(G)$) holds when G is the Rado graph. As $(G_n)_{n\in\mathbb{N}}$ is elementarily convergent to the Rado graph, this sentence holds for all but finitely many graphs G_n. Thus for all but finitely many G_n it holds $\langle \phi, G_n\rangle = \langle \widetilde{\phi}, G_n\rangle$. Moreover, according to Lemma 2.25, the sequence $(G_n)_{n\in\mathbb{N}}$ is FO-convergent and thus it holds

$$\lim_{n\to\infty} \langle \phi, G_n\rangle = \lim_{n\to\infty} \langle \widetilde{\phi}, G_n\rangle.$$

By using inclusion/exclusion argument and the general form of the density of homomorphisms of fixed target graphs to a graphon we deduce that there exists a polynomial $P_\phi \in \mathbb{Z}[X_1, \dots, X_{\binom{p}{2}}]$ (which depends only on ϕ) such that

$$\lim_{n\to\infty} \langle \widetilde{\phi}, G_n\rangle = \int \cdots \int P_\phi((W_{i,j}(x_i, x_j))_{1\leq i<j\leq p})\, \mathrm{d}x_1 \dots \mathrm{d}x_p.$$

The theorem follows. $\square$

Although elementary convergence to Rado graph seems quite a natural assumption for graphs which are neither too sparse nor too dense, elementary convergence to other ultra-homogeneous graphs may be problematic.

EXAMPLE 2.33. Cherlin [18] posed the problem whether there is a finite k-saturated triangle-free graph, for each $k \in \mathbb{N}$, where a triangle free graph is called *k-saturated* if for every set S of at most k vertices, and for every independent subset T of S, there exists a vertex adjacent to each vertex of T and to no vertex of $S - T$. In other words, Cherlin asks whether the generic countable triangle-free graph has the finite model property, that is if it is an elementary limit of a sequence of finite graphs.

It is possible to extend Lemma 2.25 to sequences of graph having a non ultra-homogeneous elementary limit if we restrict FO to a smaller fragment. An example is the following:

EXAMPLE 2.34. A graph G is *IH-Homogeneous* [15] if every partial finite isomorphism extends into an endomorphism. Let PP be the fragment of FO that consists into *primitive positive* formulas, that is formulas formed using adjacency, equality, conjunctions and existential quantification only, and let BA(PP) be the minimum sub-Boolean algebra of FO containing PP.

Following the lines of Lemma 2.25 and using Theorem 2.13 and Proposition 2.1, one proves that if a sequence of graphs $(G_n)_{n\in\mathbb{N}}$ converges elementarily to some IH-homogeneous infinite countable graph then $(G_n)_{n\in\mathbb{N}}$ is BA(PP)-convergent if and only if it is QF-convergent.

2.3.5. FO-convergence of Graphs with Bounded Maximum Degree. We now consider how full FO-convergence differs to BS-convergence for sequence of graphs with maximum degree at most D. As a corollary of Theorems 2.21 and 2.15 we have:

THEOREM 2.35. *A sequence (G_n) of finite graphs with maximum degree at most d such that $\lim_{n\to\infty} |G_n| = \infty$ is* FO*-convergent if and only if it is both BS-convergent and elementarily convergent.*

2.4. Interpretation Schemes

In the process of this research we discovered the increasing role played by interpretations. They are described in this section.

2.4.1. Continuous Functions and Interpretations. Let X and Y be fragments of $\mathrm{FO}(\kappa)$ and $\mathrm{FO}(\lambda)$, respectively. Let $f : S(\mathcal{B}(X)) \to S(\mathcal{B}(Y))$. Then f is continuous if and only if for every X-convergent sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$, the sequence $(f(\mathbf{A}_n))_{n\in\mathbb{N}}$ is Y-convergent. However, relying on the topological definition of continuity, a function $f : S(\mathcal{B}(X)) \to S(\mathcal{B}(Y))$ is continuous if and only if the inverse image of an open subset of $S(\mathcal{B}(Y))$ is an open subset of $S(\mathcal{B}(X))$. In the case of Stone spaces (where clopen subsets generates the topology), we can further restrict our attention to clopen subsets: f will be continuous if the inverse image of a clopen subset is a clopen subset. In other words, f is continuous if there exists $f^* : \mathcal{B}(Y) \to \mathcal{B}(X)$, such that for every $\phi \in Y$, it holds

$$f^{-1}(K(\phi)) = K(f^*(\phi)).$$

Note that f^* will be a homomorphism from $\mathcal{B}(Y)$ to $\mathcal{B}(X)$, and that the duality between f and f^* is nothing more the duality between Stone spaces and Boolean algebras.

The above property can be sometimes restated in terms of definable sets in structures. For a fragment X of FO and a relational structure $\mathbf{A}$, a subset $F \subseteq A^p$ is X*-definable* if there exists a formula $\phi \in X$ with free variables $x_1, \dots, x_p$ such that

$$F = \Omega_\phi(\mathbf{A}) = \{(v_1, \dots, v_p) \in A^p :\ \mathbf{A} \models \phi(v_1, \dots, v_p)\}.$$

Let $\mathbf{A}$ be a κ-structure, let $\mathbf{B}$ be a λ-structure, and let $g : A^k \to B$ be surjective. Assume that there exists a function $g^* : Y \to X$ such that for every $\phi \in Y$ with free variables $x_1, \dots, x_p$ ($p \geq 0$), and every $v_{i,j} \in A$ ($1 \leq i \leq p$, $1 \leq j \leq k$) it holds

$$\begin{gathered}\mathbf{B} \models \phi(g(v_{1,1}, \dots, v_{1,k}), \dots, g(v_{p,1}, \dots, v_{p,k}))\\ \iff \\ \mathbf{A} \models g^*(\phi)(v_{1,1}, \dots, v_{1,k}, \dots, v_{p,1}, \dots, v_{p,k})\end{gathered}$$

then g^* is a homomorphism, and thus it defines a continous function from $S(\mathcal{B}(X))$ to $S(\mathcal{B}(Y))$. Note that the above formula can be restated as

$$\Omega_{g^*(\phi)}(\mathbf{A}) = \hat{g}^{-1}(\Omega_\phi(\mathbf{B})),$$

where

$$\hat{g}((v_{1,1}, \dots, v_{p,k})) = (g(v_{1,1}, \dots, v_{1,k}), \dots, g(v_{p,1}, \dots, v_{p,k})).$$

In other words, the inverse image of a Y-definable set of $\mathbf{B}$ is an X-definable set of $\mathbf{A}$.

When $X = \mathrm{FO}(\kappa)$ and $Y = \mathrm{FO}(\lambda)$, the property that the inverse image of a first-order definable set of $\mathbf{B}$ is a first-order definable set of $\mathbf{A}$ leads to the model theoretical notion of interpretation (without parameters) of $\mathbf{B}$ in $\mathbf{A}$. We recall now the formal definition of an interpretation.

Definition 2.36 (Interpretation). An *interpretation* of $\mathbf{B}$ in $\mathbf{A}$ *with parameters* (or *without parameters*, respectively) with *exponent* k is a surjective map from a subset of A^k onto B such that the inverse image of every set X definable in $\mathbf{B}$ by a first-order formula without parameters is definable in $\mathbf{A}$ by a first-order formula with parameters (or without parameters, respectively).

2.4.2. Interpretation Schemes. The main drawback of interpretations is that it only concerns two specific structures $\mathbf{A}$ and $\mathbf{B}$. However, it is frequent that intepretations naturally generalize to a familly of interpretations of λ-structures in κ-structures with the same associated homomorphism of Boolean algebras. Moreover, this homomorphism is uniquely defined by the way it transforms each relation in λ (including equality) into a formula in κ and by the formula which defines the domain of the κ-structures. This can be formalized as follows.

Definition 2.37 (Interpretation Scheme). Let κ, λ be signatures, where λ has q relational symbols $R_1, \dots, R_q$ with respective arities $r_1, \dots, r_q$.

An *interpretation scheme* I of λ-structures in κ-structures is defined by an integer k — the *exponent* of the intrepretation scheme — a formula $E \in \mathrm{FO}_{2k}(\kappa)$, a formula $\theta_0 \in \mathrm{FO}_k(\kappa)$, and a formula $\theta_i \in \mathrm{FO}_{r_i k}(\kappa)$ for each symbol $R_i \in \lambda$, such that:

- the formula E defines an equivalence relation of k-tuples;
- each formula θ_i is compatible with E, in the sense that for every $0 \leq i \leq q$ it holds
$$\bigwedge_{1 \leq j \leq r_i} E(\mathbf{x}_j, \mathbf{y}_j) \quad \vdash \quad \theta_i(\mathbf{x}_1, \dots, \mathbf{x}_{r_i}) \leftrightarrow \theta_i(\mathbf{y}_1, \dots, \mathbf{y}_{r_i}),$$
where $r_0 = 1$, boldface $\mathbf{x}_j$ and $\mathbf{y}_j$ represent k-tuples of free variables, and where $\theta_i(\mathbf{x}_1, \dots, \mathbf{x}_{r_i})$ stands for $\theta_i(x_{1,1}, \dots, x_{1,k}, \dots, x_{r_i,1}, \dots, x_{r_i,k})$.

For a κ-structure $\mathbf{A}$, we denote by $\mathsf{I}(\mathbf{A})$ the λ-structure $\mathbf{B}$ defined as follows:

- the domain B of $\mathbf{B}$ is the subset of the E-equivalence classes $[\mathbf{x}] \subseteq A^k$ of the tuples $\mathbf{x} = (x_1, \dots, x_k)$ such that $\mathbf{A} \models \theta_0(\mathbf{x})$;
- for each $1 \leq i \leq q$ and every $\mathbf{v}_1, \dots, \mathbf{v}_{s_i} \in A^{kr_i}$ such that $\mathbf{A} \models \theta_0(\mathbf{v}_j)$ (for every $1 \leq j \leq r_i$) it holds
$$\mathbf{B} \models R_i([\mathbf{v}_1], \dots, [\mathbf{v}_{r_i}]) \quad \iff \quad \mathbf{A} \models \theta_i(\mathbf{v}_1, \dots, \mathbf{v}_{r_i}).$$

¿From the standard properties of model theoretical interpretations (see, for instance [48] p. 180), we state the following: if I is an interpretation of λ-structures in κ-structures, then there exists a mapping $\tilde{\mathsf{I}} : \mathrm{FO}(\lambda) \to \mathrm{FO}(\kappa)$ (defined by means of the formulas $E, \theta_0, \dots, \theta_q$ above) such that for every $\phi \in \mathrm{FO}_p(\lambda)$, and every κ-structure $\mathbf{A}$, the following property holds (while letting $\mathbf{B} = \mathsf{I}(\mathbf{A})$ and identifying elements of B with the corresponding equivalence classes of A^k):

For every $[\mathbf{v}_1], \dots, [\mathbf{v}_p] \in B^p$ (where $\mathbf{v}_i = (v_{i,1}, \dots, v_{i,k}) \in A^k$) it holds
$$\mathbf{B} \models \phi([\mathbf{v}_1], \dots, [\mathbf{v}_p]) \quad \iff \quad \mathbf{A} \models \tilde{\mathsf{I}}(\phi)(\mathbf{v}_1, \dots, \mathbf{v}_p).$$

It directly follows from the existence of the mapping $\tilde{\mathsf{I}}$ that an interpretation scheme I of λ-structures in κ-structures defines a continuous mapping from $S(\mathcal{B}(\mathrm{FO}(\kappa)))$ to $S(\mathcal{B}(\mathrm{FO}(\lambda)))$. Thus, interpretation schemes have the following general property:

PROPOSITION 2.3. *Let I be an interpretation scheme of λ-structures in κ-structures.*

Then, if a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of finite κ-structures is FO*-convergent then the sequence $(\mathsf{I}(\mathbf{A}_n))_{n\in\mathbb{N}}$ of (finite) λ-structures is* FO*-convergent.*

We shall be mostly interested in very specific and simple types of interpretation schemes.

DEFINITION 2.38. Let κ, λ be signatures. A *basic interpretation scheme* I of λ-structures in κ-structures with *exponent* k is defined by a formula $\theta_i \in \mathrm{FO}_{kr_i}(\kappa)$ for each symbol $R_i \in \lambda$ with arity r_i.

For a κ-structure $\mathbf{A}$, we denote by $\mathsf{I}(\mathbf{A})$ the structure with domain A^k such that, for every $R_i \in \lambda$ with arity r_i and every $\mathbf{v}_1, \dots, \mathbf{v}_{r_i} \in A^k$ it holds

$$\mathsf{I}(\mathbf{A}) \models R_i(\mathbf{v}_1, \dots, \mathbf{v}_{r_i}) \quad \iff \quad \mathbf{A} \models \theta_i(\mathbf{v}_1, \dots, \mathbf{v}_{r_i}).$$

It is immediate that every basic interpretation scheme I defines a mapping $\tilde{\mathsf{I}} : \mathrm{FO}(\lambda) \to \mathrm{FO}(\kappa)$ such that for every κ-structure $\mathbf{A}$, every $\phi \in \mathrm{FO}_p(\lambda)$, and every $\mathbf{v}_1, \dots, \mathbf{v}_p \in A^k$ it holds

$$\mathsf{I}(\mathbf{A}) \models \phi(\mathbf{v}_1, \dots, \mathbf{v}_p) \quad \iff \quad \mathbf{A} \models \tilde{\mathsf{I}}(\phi)(\mathbf{v}_1, \dots, \mathbf{v}_p)$$

and

$$\mathrm{qrank}(\tilde{\mathsf{I}}(\phi)) \le k(\mathrm{qrank}(\phi) + \max_i \mathrm{qrank}(\theta_i)).$$

It follows that for every κ-structure $\mathbf{A}$, every $\phi \in \mathrm{FO}_p(\lambda)$, it holds

$$\Omega_\phi(\mathsf{I}(\mathbf{A})) = \Omega_{\tilde{\mathsf{I}}(\phi)}(\mathbf{A}).$$

In particular, if $\mathbf{A}$ is a finite structure, it holds

$$\langle \phi, \mathsf{I}(\mathbf{A}) \rangle = \langle \tilde{\mathsf{I}}(\phi), \mathbf{A} \rangle.$$

CHAPTER 3

Modelings for Sparse Structures

3.1. Relational Samples Spaces

The notion of *relational sample space* is a strenghtening of the one of relational structure, where it is required that the domain shall be endowed with a suitable structure of (nice) measurable space.

3.1.1. Definition and Basic Properties.

DEFINITION 3.1. Let λ be a signature. A λ-*relational sample space* is a λ-structure $\mathbf{A}$, whose domain A is a standard Borel space with the property that every first-order definable subset of A^p is measurable. Precisely, for every integer p, and every $\phi \in \mathrm{FO}_p(\lambda)$, denoting

$$\Omega_\phi(\mathbf{A}) = \{(v_1, \dots, v_p) \in A^p : \ \mathbf{A} \models \phi(v_1, \dots, v_p)\},$$

it holds $\Omega_\phi(\mathbf{A}) \in \Sigma_{\mathbf{A}}^p$, where $\Sigma_{\mathbf{A}}$ is the Borel σ-algebra of A.

Note, that in the case of graphs, every relational sample space is a Borel graph (that is a graph whose vertex set is a standard Borel space and whose edge set is Borel), but the converse is not true.

LEMMA 3.2. *Let λ be a signature, let $\mathbf{A}$ be a λ-structure, whose domain A is a standard Borel space with σ-algebra $\Sigma_{\mathbf{A}}$.*

Then the following conditions are equivalent:

(a) $\mathbf{A}$ is a λ-relational sample space;
(b) for every integer $p \geq 0$ and every $\phi \in \mathrm{FO}_p(\lambda)$, it holds $\Omega_\phi(\mathbf{A}) \in \Sigma_{\mathbf{A}}^p$;
(c) for every integer $p \geq 1$ and every $\phi \in \mathrm{FO}_p^{\mathrm{local}}(\lambda)$, it holds $\Omega_\phi(\mathbf{A}) \in \Sigma_{\mathbf{A}}^p$;
(d) for every integers $p, q \geq 0$, every $\phi \in \mathrm{FO}_{p+q}(\lambda)$, and every $a_1, \dots, a_q \in A^q$ the set

$$\{(v_1, \dots, v_p) \in A^p : \ \mathbf{A} \models \phi(a_1, \dots, a_q, v_1, \dots, v_p)\}$$

belongs to $\Sigma_{\mathbf{A}}^p$.

PROOF. Items (a) and (b) are equivalent by definition. Also we obviously have the implications (d) $\Rightarrow$ (b) $\Rightarrow$ (c). That (c) $\Rightarrow$ (b) is a direct consequence of Gaifman locality theorem, and the implication (b) $\Rightarrow$ (d) is a direct consequence of Fubini-Tonelli theorem. □

LEMMA 3.3. *Let $\mathbf{A}$ be a relational sample space, let $a \in A$, and let $\mathbf{A}_a$ be the connected component of $\mathbf{A}$ containing a.*

Then $\mathbf{A}_a$ has a measurable domain and, equipped with the σ-algebra of the Borel sets of A included in A_a, it is a relational sample space.

PROOF. Let $\phi \in \mathrm{FO}_p^{\mathrm{local}}$ and let

$$X = \{(v_1, \dots, v_p) \in A_a^p : \ \mathbf{A}_a \models \phi(v_1, \dots, v_p)\}.$$

As ϕ is local, there is an integer D such that the satisfaction of ϕ only depends on the D-neighborhoods of the free variables.

For every integer $n \in \mathbb{N}$, denote by $B(\mathbf{A}, a, n)$ the substructure of $\mathbf{A}$ induced by all vertices at distance at most n from a. By the locality of ϕ, for every $v_1, \dots, v_p$ at distance at most n from a it holds

$$\mathbf{A}_a \models \phi(v_1, \dots, v_p) \qquad \iff \qquad B(\mathbf{A}, a, n + D) \models \phi(v_1, \dots, v_p).$$

However, it is easily checked that there is a local first-order formula $\varphi_n \in \mathrm{FO}_{p+1}^{\mathrm{local}}$ such that for every $v_1, \dots, v_p$ it holds

$$B(\mathbf{A}, a, n{+}D) \models \phi(v_1, \dots, v_p) \wedge \bigwedge_{i=1}^{p} \mathrm{dist}(a, v_i) \leq n \qquad \iff \qquad \mathbf{A} \models \varphi_n(a, v_1, \dots, v_p).$$

By Lemma 3.2, it follows that the set $X_n = \{(v_1, \dots, v_n) \in A : \ \mathbf{A} \models \varphi_n(a, v_1, \dots, v_p)\}$ is measurable. As $X = \bigcup_{n \in \mathbb{N}} X_n$, we deduce that X is measurable (with respect to $\Sigma_{\mathbf{A}}^p$. In particular, A_a is a Borel subset of A hence A_a, equipped with the σ-algebra $\Sigma_{\mathbf{A}_a}$ of the Borel sets of A included in A_a, is a standard Borel set. Moreover, it is immediate that a subset of A_a^p belongs to $\Sigma_{\mathbf{A}_a}^p$ if and only if it belongs to $\Sigma_{\mathbf{A}}^p$. Hence, every subset of A_a^p defined by a local formula is measurable with respect to $\Sigma_{\mathbf{A}_a}^p$. By Lemma 3.2, it follows that $\mathbf{A}_a$ is a relational sample space. $\square$

3.1.2. Interpretions of Relational Sample Spaces. An elementary interpretation with parameter consits in distinguishing a single element, the parameter, by adding a new unary symbol to the signature (e.g. representing a root).

LEMMA 3.4. *Let* $\mathbf{A}$ *be a* λ*-relational sample space, let* λ^+ *be the signature obtained from* λ *by adding a new unary symbol* M *and let* $\mathbf{A}^+$ *be obtained from* $\mathbf{A}$ *by marking a single* $a \in A$ *(i.e.* a *is the only element* x *of* $A^+ = A$ *such that* $\mathbf{A}^+ \models M(x)$*).*

Then $\mathbf{A}^+$ *is a relational sample space.*

PROOF. Let $\phi \in \mathrm{FO}_p(\lambda^+)$. There exists $\phi' \in \mathrm{FO}_{p+1}(\lambda)$ such that for every $x_1, \dots, x_p \in A$ it holds

$$\mathbf{A}^+ \models \phi(x_1, \dots, x_p) \qquad \iff \qquad \mathbf{A} \models \phi(a, x_1, \dots, x_p).$$

According to Lemma 3.2, the set of all $(x_1, \dots, x_p)$ such that $\mathbf{A} \models \phi(a, x_1, \dots, x_p)$ is measurable. It follows that $\mathbf{A}^+$ is a relational sample space. $\square$

LEMMA 3.5. *Every injective first-order interpretation (with or without parameters) of a relational sample space is a relational sample space.*

Precisely, if f *is an injective first-order interpretation of a* λ*-structure* $\mathbf{B}$ *in a* κ*-relational sample space* $\mathbf{A}$ *and if we define*

$$\Sigma_{\mathbf{B}} = \{X \subseteq B : f^{-1}(X) \in \Sigma_{\mathbf{A}}^k\},$$

then $(\mathbf{B}, \Sigma_{\mathbf{B}})$ *is a relational sample space.*

Proof. According to Lemma 3.4, we can first mark all the parameters and reduce to the case were the interpretation has no parameters.

Let D be the domain of f. As B is first-order definable in $\mathbf{B}$, D is first-order definable in $\mathbf{A}$ hence $D \in \Sigma_{\mathbf{A}}^k$. Then D is a Borel sub-space of A^k. As f is a bijection from D to B, we deduce that $(B, \Sigma_{\mathbf{B}})$ is a standard Borel space.

Moreover, as the inverse image of every first-order definable set of $\mathbf{B}$ is first-order definable in $\mathbf{A}$, we deduce that $(\mathbf{B}, \Sigma_{\mathbf{B}})$ is a λ-relational sample space. □

3.1.3. Disjoint union. Let $\mathbf{H}_i$ be λ-relational sample spaces for $i \in I \subseteq \mathbb{N}$. We define the *disjoint union*

$$\mathbf{H} = \coprod_{i\in I} \mathbf{H}_i$$

of the $\mathbf{H}_i$'s as the relational structure, which is the disjoint union of the $\mathbf{H}_i$'s endowed with the σ-algebra $\Sigma_{\mathbf{H}} = \{\bigcup_i X_i : \ X_i \in \Sigma_{\mathbf{H}_i}\}$.

Lemma 3.6. *Let $\mathbf{H}_i$ be λ-relational sample spaces for $i \in I \subseteq \mathbb{N}$. Then $\mathbf{H} = \coprod_{i\in I} \mathbf{H}_i$ is a λ-relational sample space, in which every H_i is measurable.*

Proof. We consider the signature λ^+ obtained from λ by adding a new binary relation ϖ, and the basic interpretation scheme I_1 of λ^+-structures in λ-structures corresponding to the addition of the new relation ϖ by the formula $\theta_\varpi = 1$. This means that for every λ-structure $\mathbf{A}$ it holds $\mathsf{I}_1(\mathbf{A}) \models (\forall x, y)\ \varpi(x, y)$. Let $\mathbf{H}_i^+ = \mathsf{I}_1(\mathbf{H}_i)$.

Let $\mathbf{H}^+ = \coprod_{i\in I} \mathbf{H}_i^+$. Clearly, $\Sigma_{\mathbf{H}^+} = \Sigma_{\mathbf{H}}$ and $(H, \Sigma_{\mathbf{H}})$ is a standard Borel space. Moreover, by construction, each H_i is measurable.

Let $\phi \in \mathrm{FO}_p(\lambda)$. First notice that for every $(v_1, \dots, v_p) \in H^{p+q}$ (which is also $(H^+)^{p+q}$) it holds $\Omega_\phi(\mathbf{H}) = \Omega_\phi(\mathbf{H}^+)$, that is:

$$\mathbf{H} \models \phi(v_1, \dots, v_p) \quad\iff\quad \mathbf{H}^+ \models \phi(v_1, \dots, v_p).$$

It follows from Lemma 2.23 that the set $\Omega_\phi(\mathbf{H}^+)$ may be obtained by Boolean operations, products, and coordinate permuations from sets defined by ϖ-local formulas (which we introduced in Section 2.3.3). As all these operations preserve measurability, we can assume that ϕ is ϖ-local. Then $\Omega_\phi(\mathbf{H}^+)$ is the union of the sets $\Omega_\phi(\mathbf{H}_i)$. All these sets are measurable (as $\mathbf{H}_i$ is a modeling) thus their union is measurable (by construction of $\Sigma_{\mathbf{H}}$). It follows that $\mathbf{H}^+$ is a relational sample space, and so is $\mathbf{H}$ (every first-order definable set of $\mathbf{H}$ is first-order definable in $\mathbf{H}^+$). □

3.2. Modelings

We introduced a notion of limit objects — called modelings — for sequences of sparse graphs and structures, which is a natural generalization of graphings. These limit objects are defined by considering a probability measure on a relational sample space. In this section, we show that the most we can expect is that modelings are limit objects for sequence of sparse structures, and we conjecture that an unavoidable qualtitative jump occurs for notions of limit structures, which coincides with the nowhere dense/somewhere dense fronteer (see Conjecture 1.1).

3.2.1. Definition and Basic Properties. Recall Definitions 1.4 and 1.6: a λ-*modeling* $\mathbf{A}$ is a λ-relational sample space equipped with a probability measure (denoted $\nu_{\mathbf{A}}$), and the *Stone pairing* of $\phi \in \mathrm{FO}(\lambda)$ and a λ-modeling $\mathbf{A}$ is $\langle \phi, \mathbf{A}\rangle = \nu_{\mathbf{A}}^p(\Omega_\phi(\mathbf{A}))$. Notice that it follows (by Fubini's theorem) that it holds

$$\begin{aligned}\langle \phi, \mathbf{A}\rangle &= \int_{\mathbf{x}\in A^p} 1_{\Omega_\phi(\mathbf{A})}(\mathbf{x})\,\mathrm{d}\nu_{\mathbf{A}}^p(\mathbf{x})\\ &= \int \cdots \int \mathbf{1}_{\Omega_\phi(\mathbf{A})}(x_1,\dots,x_p)\,\mathrm{d}\nu_{\mathbf{A}}(x_1)\,\dots\,\mathrm{d}\nu_{\mathbf{A}}(x_p).\end{aligned}$$

Then, generalizing Definition 1.7, we extend the notion of X-convergence to modelings:

DEFINITION 3.7 (modeling X-limit). Let X be a fragment of $\mathrm{FO}(\lambda)$. If an X-convergent sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of λ-modelings satisfies

$$(\forall \phi \in X) \quad \langle \phi, \mathbf{L}\rangle = \lim_{n\to\infty} \langle \phi, \mathbf{A}_n\rangle$$

for some λ-modeling $\mathbf{L}$, then we say that $\mathbf{L}$ is a *modeling* X*-limit* of $(\mathbf{A}_n)_{n\in\mathbb{N}}$.

A λ-modeling $\mathbf{A}$ is *weakly uniform* if all the singletons of A have the same measure. Clearly, every finite λ-structure $\mathbf{A}$ can be identified with the weakly uniform modeling obtained by considering the discrete topology on A. This identification is clearly consistent with our definition of the Stone pairing of a formula and a modeling.

In the case where a modeling $\mathbf{A}$ has an infinite domain, the condition for $\mathbf{A}$ to be weakly uniform is equivalent to the condition for $\nu_{\mathbf{A}}$ to be atomless. This property is usually fulfilled by modeling X-limits of sequences of finite structures.

LEMMA 3.8. *Let X be a fragment of* FO *that includes* FO_0 *and the formula* $(x_1 = x_2)$. *Then every modeling X-limit of weakly uniform modelings is weakly uniform.*

PROOF. Let ϕ be the formula $(x_1 = x_2)$. Notice that for every finite λ-structure $\mathbf{A}$ it holds $\langle \phi, \mathbf{A}\rangle = 1/|A|$ and that for every infinite weakly uniform λ-structure it holds $\langle \phi, \mathbf{A}\rangle = 0$.

Let $\mathbf{L}$ be a modeling X-limit of a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$. Assume $\lim_{n\to\infty} |A_n| = \infty$. Assume for contradiction that $\nu_{\mathbf{L}}$ has an atom $\{v\}$ (i.e. $\nu_{\mathbf{L}}(\{v\}) > 0$). Then $\langle \phi, \mathbf{L}\rangle \geq \nu_{\mathbf{L}}(\{v\})^2 > 0$, contradicting $\lim_{n\to\infty}\langle \phi, \mathbf{A}_n\rangle = 0$. Hence $\nu_{\mathbf{L}}$ is atomless.

Otherwise, $|L| = \lim_{n\to\infty} |A_n| < \infty$ (as $\mathbf{L}$ is an elementary limit of $(\mathbf{A}_n)_{n\in\mathbb{N}}$). Let $N = |L|$. Label $v_1,\dots,v_N$ the elements of L and let $p_i = \nu_{\mathbf{L}}(\{v_i\})$. Then

$$\begin{aligned}\frac{1}{N}\sum_{i=1}^N p_i^2 - \Big(\frac{1}{N}\sum_{i=1}^N p_i\Big)^2 &= \frac{\langle \phi, \mathbf{L}\rangle}{N} - \frac{1}{N^2}\\ &= \frac{\lim_{n\to\infty}\langle \phi, \mathbf{A}_n\rangle}{N} - \frac{1}{N^2}\\ &= 0\end{aligned}$$

Thus $p_i = 1/N$ for every $i = 1,\dots,N$. $\square$

COROLLARY 3.1. *Every modeling* $\mathrm{FO}_2^{\mathrm{local}}$*-limit of finite structures is weakly uniform.*

LEMMA 3.9. *Let X be a fragment that includes all quantifier free formulas.*

Assume $\mathbf{L}$ is a modeling X-limit of a sequence $(G_n)_{n\in\mathbb{N}}$ of graphs with $|G_n|\to\infty$. Let $\overline{\nu}_{\mathbf{L}}$ be the completion of the measure $\nu_{\mathbf{L}}$.

Then there is at least one mod 0 isomorphism $f:[0,1]\to(L,\overline{\nu}_{\mathbf{L}})$, and for every such f the graphon W defined by

$$W(x,y)=\mathbf{1}_{\Omega_{(x_1\sim x_2)}(\mathbf{L})}(f(x),f(y))$$

(for x,y in the domain of f, and $W(x,y)=0$ elsewhere) is as a random-free graphon L-limit of $(G_n)_{n\in\mathbb{N}}$.

PROOF. Considering the formula $x_1=x_2$, we infer that $\nu_{\mathbf{L}}$ is atomless. This measure is also atomless and turns L into a standard probability space. According to the isomorphism theorem, all atomless standard probability spaces are mutually mod 0 isomorphic hence there is at least one mod 0 isomorphism $f:[0,1]\to(L,\overline{\nu}_{\mathbf{L}})$ ($[0,1]$ is considered with Lebesgue measure).

Fix such a mod 0 isomorphism f, defined on $[0,1]\setminus N_1$, with value on $L\setminus N_2$ (where N_1 and N_2 are nullsets). For every Borel measurable function $g:L^n\to[0,1]$, define g_f by $g_f(x_1,\dots,x_n)=g(f(x_1),\dots,f(x_n))$ if $x_i\notin N_1$ for every $1\leq i\leq n$ and $g_f(x_1,\dots,x_n)=0$ otherwise. Then it holds

$$\begin{aligned}\int_{[0,1]^n} g_f(x_1,\dots,x_n)\,\mathrm{d}x_1\dots\mathrm{d}x_n &= \int_{L^n} g(v_1,\dots,v_n)\,\mathrm{d}\overline{\nu}_{\mathbf{L}}(v_1)\dots\mathrm{d}\overline{\nu}_{\mathbf{L}}(v_n)\\ &= \int_{L^n} g(v_1,\dots,v_n)\,\mathrm{d}\nu_{\mathbf{L}}(v_1)\dots\mathrm{d}\nu_{\mathbf{L}}(v_n).\end{aligned}$$

It follows that for every finite graph F with vertex set $\{1,\dots,n\}$, denoting ϕ_F the formula $\bigwedge_{ij\in E(F)}(x_i\sim x_j)$, it holds

$$\begin{aligned}t(F,W) &= \int_{[0,1]^n}\prod_{ij\in E(F)} W(x_i,x_j)\,\mathrm{d}x_1\dots\mathrm{d}x_n\\ &= \int_{L^n}\prod_{ij\in E(F)}\mathbf{1}_{\Omega_{(x_1\sim x_2)}(\mathbf{L})}(v_i,v_j)\,\mathrm{d}\nu_{\mathbf{L}}(v_1)\dots\mathrm{d}\nu_{\mathbf{L}}(v_n)\\ &= \int_{L^n}\mathbf{1}_{\Omega_{\phi_F}(\mathbf{L})}(v_1,\dots,v_n)\,\mathrm{d}\nu_{\mathbf{L}}(v_1)\dots\mathrm{d}\nu_{\mathbf{L}}(v_n)\\ &= \langle\phi_F,\mathbf{L}\rangle\\ &= \lim_{n\to\infty}\langle\phi_F,G_n\rangle\\ &= \lim_{n\to\infty} t(F,G_n).\end{aligned}$$

Hence W is a graphon L-limit of $(G_n)_{n\in\mathbb{N}}$. As W is $\{0,1\}$-valued, it is (by definition) random-free. $\square$

We deduce the following limitation of modelings as limit objects.

COROLLARY 3.2. *Let X be a fragment that includes all quantifier free formulas.*

Assume $(G_n)_{n\in\mathbb{N}}$ is an X-convergent sequence of graphs with unbounded order, which is L-convergent to some non random-free graphon W. Then $(G_n)_{n\in\mathbb{N}}$ has no modeling X-limit.

Let us now give some example stressing that the nullsets of the mod 0 isomorphism f can be quite large, making $\mathbf{L}$ and W look quite different. We give now an example in the more general setting of directed graphs and non-symmetric graphons.

Example 3.10. Let $\vec{T}_n$ be the transitive tournament of order n, that is the directed graph on $\{1,\dots,n\}$ defined from the natural linear order $<_n$ on $\{1,\dots,n\}$ by $i \to j$ if $i < j$. This sequence is obviously FO-convergent.

It is not difficult to construct a modeling FO-limit of $(\vec{T}_n)_{n\in\mathbb{N}}$: Let

$$L = \{0\} \times \mathbb{Z}^+ \cup]0,1[\times \mathbb{Z} \cup \{1\} \times \mathbb{Z}^-,$$

with the Borel σ-algebra Σ generated by the product topology of $\mathbb{Z}$ (with discrete topology) and $\mathbb{R}$ (with usual topology). On L we define a linear order $<_L$ by $(\alpha,i) <_L (\beta,j)$ if $\alpha < \beta$ or $(\alpha = \beta)$ and $(i < j)$. That (L,Σ) is a relational sample space follows from the o-minimality of $([0,1],<)$. The measure $\nu_{\mathbf{L}}$ can be defined as the product of Lebesgue measure on $[0,1]$ by any probability measure on $\mathbb{Z}$. For instance, for every $B \in \Sigma$ we let $\nu_{\mathbf{L}}(B) = \lambda(B \cap ([0,1] \times \{0\}))$, where λ is Lebesgue measure. It is not difficult to check that $\mathbf{L}$ is indeed a modeling FO-limit of $(\{1,\dots,n\},<_n) \simeq \vec{T}_n$.

In this case, a mod 0 isomorphism $f : [0,1] \to L$ can be defined by $f(x) = (x,0)$. The null set N_2, although very large, is clearly a $\nu_{\mathbf{L}}$-nullset, and the obtained (non symmetric) random-free graphon $W : [0,1] \times [0,1] \to [0,1]$ is simply defined by

$$W(x,y) = \begin{cases} 1 & \text{if } x < y \\ 0 & \text{otherwise.} \end{cases}$$

Note that W corresponds to $[0,1]$ with its natural order $<$. This order is clearly an L-limit of $<_n$ (but not an elementary limit, as it is dense although no finite order is).

In the spirit of Lemma 3.2, we propose the following problems:

Problem 3.1. Let $\mathbf{L}$ be a modeling FO-limit of a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of λ-structures, and let $v \in L$. Does there exist a sequence $(v_n)_{n\in\mathbb{N}}$ such that $v_n \in A_n$ and such that the rooted modeling $(\mathbf{L}, v)$ is a modeling FO-limit of the rooted structures $(\mathbf{A}_n, v_n)$?

Problem 3.2. Let $\mathbf{L}$ be a modeling FO-limit of a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of λ-structures. Does there exist $f : L \to \prod_{i\in\mathbb{N}} A_n$ such that for every $v_1,\dots,v_k \in L$, the k-rooted modeling $(\mathbf{L}, v_1,\dots,v_k)$ is a modeling FO-limit of the k-rooted structures $(\mathbf{A}_n, f(v_1)_n, \dots, f(v_k)_n)$?

3.2.2. Interpretation Schemes applied to Modelings. Basic interpretation schemes will be an efficient tool to handle modelings. Let I be an interpretation scheme I of λ-structures in κ-structures. We have seen that I can be extended to a mapping from κ-relational sample space to λ-relational sample space. In the case where I is a basic interpretation scheme, we further extend I to a mapping from κ-modeling to λ-modeling: For a κ-modeling $\mathbf{A}$, the λ-modeling $\mathbf{B} = \mathsf{I}(\mathbf{A})$ is the modeling on the image relational sample space of $\mathbf{A}$ with the probability measure $\nu_{\mathbf{B}} = \nu_{\mathbf{A}}$. This is formalized as follows:

Lemma 3.11. *Let I be a basic interpretation scheme I of λ-structures in κ-structures with exponent k. Extend the definition of I to a mapping of κ-modeling* *to λ-modeling by setting $\nu_{\mathsf{I}(\mathbf{A})} = \nu^k_{\mathbf{A}}$. Then for every κ-modeling $\mathbf{A}$ and every $\phi \in \mathrm{FO}(\lambda)$ it holds*

$$\langle \phi, \mathsf{I}(\mathbf{A})\rangle = \langle \widetilde{\mathsf{I}}(\phi), \mathbf{A}\rangle.$$

PROOF. Let $\mathbf{A}$ be a κ-modeling. For every $\phi \in \mathrm{FO}_p(\lambda)$ it holds

$$\Omega_\phi(\mathsf{I}(\mathbf{A})) = \Omega_{\widetilde{\mathsf{I}}(\phi)}(\mathbf{A})$$

thus $\langle \phi, \mathsf{I}(\mathbf{A})\rangle = \nu^p_{\mathsf{I}(\mathbf{A})}(\Omega_\phi(\mathsf{I}(\mathbf{A}))) = \nu^{kp}_{\mathbf{A}}(\Omega_{\widetilde{\mathsf{I}}(\phi)}(\mathbf{A})) = \langle \widetilde{\mathsf{I}}(\phi), \mathbf{A}\rangle$. □

REMARK 3.12. If the basic interpretation scheme I is defined by quantifier-free formulas only, then it is possible to define $\widetilde{\mathsf{I}}$ in such a way that for every $\phi \in \mathrm{FO}(\lambda)$ it holds $\mathrm{qrank}(\widetilde{\mathsf{I}}(\phi)) \leq \mathrm{qrank}(\phi)$.

The following strengthening of Proposition 2.3 in the case where we consider a basic interpretation scheme is a clear consequence of Lemma 3.11.

PROPOSITION 3.1. *Let I be a basic interpretation scheme of λ-structures in κ-structures.*

If $\mathbf{L}$ is a modeling FO-limit of a sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ of κ-modelings then $\mathsf{I}(\mathbf{L})$ is a modeling FO-limit of the sequence $(\mathsf{I}(\mathbf{A}_n))_{n\in\mathbb{N}}$.

LEMMA 3.13. *Let $p \in \mathbb{N}$ be a positive integer, let $\mathbf{L}$ be a modeling, and let ${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp} : L^p \to S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ be the function mapping $(v_1, \dots, v_p) \in L^p$ to the complete theory of $(\mathbf{L}, v_1, \dots, v_p)$ (that is the set of the formulas $\varphi \in \mathrm{FO}_p(\lambda)$ such that $\mathbf{L} \models \varphi(v_1, \dots, v_p)$).*

Then ${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}$ is a measurable map from $(L^p, \Sigma^p_{\mathbf{L}})$ to $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ (with its Borel σ-algebra).
Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be an $\mathrm{FO}_p(\lambda)$-convergent sequence of finite λ-structures, and let μ_p be the associated limit measure (as in Theorem 2.6).

Then $\mathbf{L}$ is an $\mathrm{FO}_p(\lambda)$-limit modeling of $(\mathbf{A}_n)_{n\in\mathbb{N}}$ if and only if μ_p is the push-forward of the product measure $\nu^p_{\mathbf{L}}$ by the measurable map ${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}$, that is:

$${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}_*(\nu^p_{\mathbf{L}}) = \mu_p.$$

PROOF. Recall that the clopen sets of $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ are of the form $K(\phi)$ for $\phi \in \mathrm{FO}_p(\lambda)$ and that they generate the topology of $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ hence also its Borel σ-algebra.

That ${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}$ is measurable follows from the fact that for every $\phi \in \mathrm{FO}_p$ the preimage of $K(\phi)$, that is ${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}^{-1}(K(\phi)) = \Omega_\phi(\mathbf{L})$, is measurable.

Assume that $\mathbf{L}$ is an $\mathrm{FO}_p(\lambda)$-limit modeling of $(\mathbf{A}_n)_{n\in\mathbb{N}}$. In order to prove that ${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}_*(\nu^p_{\mathbf{L}}) = \mu_p$, it is sufficient to check it on sets $K(\phi)$:

$$\mu_p(K(\phi)) = \lim_{n\to\infty} \langle \phi, \mathbf{A}_n\rangle = \langle \phi, \mathbf{L}\rangle = \nu^p_{\mathbf{L}}({}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}^{-1}(K(\phi))).$$

Conversely, if ${}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}_*(\nu^p_{\mathbf{L}}) = \mu_p$ then for every $\phi \in \mathrm{FO}_p(\lambda)$ it holds

$$\langle \phi, \mathbf{L}\rangle = \nu^p_{\mathbf{L}}({}^{\mathsf{p}}_{\mathsf{L}}\mathsf{Tp}^{-1}(K(\phi))) = \mu_p(K(\phi)) = \lim_{n\to\infty}\langle \phi, \mathbf{A}_n\rangle,$$

hence $\mathbf{L}$ is an $\mathrm{FO}_p(\lambda)$-limit modeling of $(\mathbf{A}_n)_{n\in\mathbb{N}}$. □

If (X, Σ) is a Borel space with a probability measure ν, it is standard to define the product σ-algebra Σ^ω on the infinite product space $X^{\mathbb{N}}$, which is generated by cylinder sets of the form

$$R = \{f \in L^{\mathbb{N}} :\ f(i_1) \in A_{i_1}, \dots, f(i_k) \in A_{i_k}\}$$

for some $k \in \mathbb{N}$ and $A_{i_1}, \dots, A_{i_k} \in \Sigma$. The measure ν^ω of the cylinder R defined above is then

$$\nu^\omega(R) = \prod_{j=1}^{k} \nu(A_{i_j}).$$

By Kolmogorov's Extension Theorem, this extends to a unique probability measure on Σ^ω (which we still denote by ν^ω). We summarize this as the following (see also Fig. 3.2.2).

THEOREM 3.14. *let* $\mathbf{L}$ *be a modeling, and let* ${}^{\omega}_{\mathbf{L}}\mathsf{Tp} : L^{\mathbb{N}} \to S(\mathcal{B}(\mathrm{FO}(\lambda)))$ *be the function which assigns to* $f \in L^{\mathbb{N}}$ *the point of* $S(\mathcal{B}(\mathrm{FO}(\lambda)))$ *corresponding to the set* $\{\phi : \mathbf{L} \models \phi(f(1), \dots, f(p)),\ \textit{where}\ \mathrm{Fv}(\phi) \subseteq \{1, \dots, p\}\}$.

Then ${}^{\omega}_{\mathbf{L}}\mathsf{Tp}$ *is a measurable map.*

Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *be an* $\mathrm{FO}(\lambda)$*-convergent sequence of finite* λ*-structures, and let* μ *be the associated limit measure (see Theorem 2.4).*

Then $\mathbf{L}$ *is an* $\mathrm{FO}(\lambda)$*-limit modeling of* $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *if and only if*

$${}^{\omega}_{\mathbf{L}}\mathsf{Tp}_*(\nu^{\omega}_{\mathbf{L}}) = \mu.$$

Fig. 3.2.2 visualizes Lemma 3.13 and Theorem 3.14.

$$\begin{tikzcd}
 & \begin{array}{c} L \\ \text{with } \nu_{\mathbf{L}} \end{array} \arrow[d, "{}^{1}_{\mathbf{L}}\mathsf{Tp}"'] & \begin{array}{c} L^2 \\ \text{with } \nu^2_{\mathbf{L}} \end{array} \arrow[l] \arrow[loop above, "S_2"] \arrow[d, "{}^{2}_{\mathbf{L}}\mathsf{Tp}"'] & \cdots \arrow[l] & \begin{array}{c} L^p \\ \text{with } \nu^p_{\mathbf{L}} \end{array} \arrow[l] \arrow[loop above, "S_p"] \arrow[d, "{}^{p}_{\mathbf{L}}\mathsf{Tp}"'] & \begin{array}{c} L^{\mathbb{N}} \\ \text{with } \nu^\omega_{\mathbf{L}} \end{array} \arrow[l] \arrow[loop above, "S_\omega"] \arrow[d, "{}^{\omega}_{\mathbf{L}}\mathsf{Tp}"'] \\
\begin{array}{c} S(\mathcal{B}(\mathrm{FO}_0)) \\ \text{with } \delta_T \end{array} & \begin{array}{c} S(\mathcal{B}(\mathrm{FO}_1)) \\ \text{with } \mu_1 \end{array} \arrow[l] & \begin{array}{c} S(\mathcal{B}(\mathrm{FO}_2)) \\ \text{with } \mu_2 \end{array} \arrow[l] \arrow[loop above, "S_2"] \arrow[d] & \cdots \arrow[l] & \begin{array}{c} S(\mathcal{B}(\mathrm{FO}_p)) \\ \text{with } \mu_p \end{array} \arrow[l] \arrow[loop above, "S_p"] \arrow[d] & \begin{array}{c} S(\mathcal{B}(\mathrm{FO})) \\ \text{with } \mu \end{array} \arrow[l] \arrow[loop above, "S_\omega"] \arrow[d] \\
 & & \begin{array}{c} S(\mathcal{B}(\mathrm{FO}_1))^2 \\ \text{with } \mu_1^2 \end{array} \arrow[ul] \arrow[loop above, "S_2"] & \cdots \arrow[l] & \begin{array}{c} S(\mathcal{B}(\mathrm{FO}_1))^p \\ \text{with } \mu_1^p \end{array} \arrow[l] \arrow[loop above, "S_p"] & \begin{array}{c} S(\mathcal{B}(\mathrm{FO}_1))^{\mathbb{N}} \\ \text{with } \mu_1^\omega \end{array} \arrow[l] \arrow[loop above, "S_\omega"]
\end{tikzcd}$$

FIGURE 1. Pushforward of measures

REMARK 3.15. We could have considered that free variables are indexed by $\mathbb{Z}$ instead of $\mathbb{N}$. In such a context, natural shift operations S and T act respectively on the Stone space $\mathcal{S}$ of the Lindendaum-Tarski algebra of $\mathrm{FO}(\lambda)$, and on the space $\mathbf{L}^{\mathbb{Z}}$ of the mappings from $\mathbb{Z}$ to a λ-modeling $\mathbf{L}$. If $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is an FO-convergent sequence with limit measure μ on $\mathcal{S}$, then $(\mathcal{S}, \mu, S)$ is a measure-preserving dynamical system. Also, if $\nu^{\mathbb{Z}}$ is the product measure on $\mathbf{A}$, $(\mathbf{A}^{\mathbb{Z}}, \nu, T)$ is a Bernoulli scheme. Then, the condition of Theorem 3.14 can be restated as follows: the modeling $\mathbf{L}$ is a modeling FO-limit of the sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ if and only if $(\mathcal{S},\mu,S)$ is a factor of $(\mathbf{A}^{\mathbb{Z}},\nu^{\mathbb{Z}},T)$. This setting leads to yet another interpretation of our result, which we hope will be treated elsewhere.

3.2.3. Component-Local Formulas. The basis observation is that for ϖ-local formulas, we can reduce the Stone pairing to components.

LEMMA 3.16. *Let $\mathbf{A}$ be a λ-modeling and component relation ϖ. Let $\psi \in \mathrm{FO}_p(\lambda)$ be a ϖ-local formula of $\mathbf{A}$.*

Assume $\mathbf{A}$ has countably many connected components $\{\mathbf{A}_i\}_{i\in\Gamma}$. Let Γ_+ be the set of indexes i such that $\nu_{\mathbf{A}}(A_i) > 0$. For $i \in \Gamma_+$ we equip $\mathbf{A}_i$ with the σ-algebra $\Sigma_{\mathbf{A}_i}$ and the probability measure $\nu_{\mathbf{A}_i}$, where $\Sigma_{\mathbf{A}_i}$ is restriction of $\Sigma_{\mathbf{A}}$ to A_i and, for $X \in \Sigma_{\mathbf{A}_i}$, $\nu_{\mathbf{A}_i}(X) = \nu_{\mathbf{A}}(X)/\nu_{\mathbf{A}}(A_i)$. Then

$$\langle \psi, \mathbf{A}\rangle = \sum_{i\in\Gamma} \nu_{\mathbf{A}}(A_i)^p \,\langle \psi, \mathbf{A}_i\rangle.$$

PROOF. First note that each connected component of $\mathbf{A}$ is measurable: let $\mathbf{A}_i$ be a connected component of $\mathbf{A}$ and let $a \in A_i$. Then $A_i = \{x \in A : \mathbf{A} \models \varpi(x,a)\}$ hence A_i is measurable as $\mathbf{A}$ is a relational sample space. Let $Y = \{(v_1,\dots,v_p) \in A^p : \mathbf{A} \models \psi(v_1,\dots,v_p)\}$. Then $\langle \psi, \mathbf{A}\rangle = \nu_{\mathbf{A}}^p(Y)$. As ψ is ϖ-local, it also holds $Y = \bigcup_{i\in\Gamma} Y_i$, where $Y_i = \{(v_1,\dots,v_p) : \mathbf{A}_i \models \psi(v_1,\dots,v_p)\} = Y \cap A_i^p$. As $A_i \in \Sigma_{\mathbf{A}}$ and $Y \in \Sigma_{\mathbf{A}}^p$, it follows that $Y_i \in \Sigma_{\mathbf{A}}^p$ and (by countable additivity) it holds

$$\langle \psi, \mathbf{A}\rangle = \nu_{\mathbf{A}}^p(Y) = \sum_{i\in\Gamma} \nu_{\mathbf{A}}^p(Y_i) = \sum_{i\in\Gamma_+} \nu_{\mathbf{A}}(\mathbf{A}_i)^p \nu_{\mathbf{A}_i}^p(Y_i) = \sum_{i\in\Gamma} \nu_{\mathbf{A}}(\mathbf{A}_i)^p \,\langle \psi, \mathbf{A}_i\rangle.$$

□

COROLLARY 3.3. *Let $\mathbf{A}$ be a finite λ-structure with component relation ϖ. Let $\psi \in \mathrm{FO}_p(\lambda)$ be a ϖ-local formula of $\mathbf{A}$.*

Let $\mathbf{A}_1,\dots,\mathbf{A}_n$ be the connected components of $\mathbf{A}$. Then

$$\langle \psi, \mathbf{A}\rangle = \sum_{i=1}^{n} \left(\frac{|A_i|}{|A|}\right)^p \langle \psi, \mathbf{A}_i\rangle.$$

We are now ready to reduce Stone pairing of local formulas to Stone pairings with ϖ-local formulas on connected components.

THEOREM 3.17. *Let $p \in \mathbb{N}$ and $\phi \in \mathrm{FO}_p^{\mathrm{local}}(\lambda)$.*

Then there exist ϖ-local formulas $\xi_{i,j} \in \mathrm{FO}_{q_{i,j}}^{\mathrm{local}}$ ($1 \le i \le n$, $j \in I_i$) with $\mathrm{qrank}(\xi_{i,j}) \le \mathrm{qrank}(\phi)$ such that for each i, $\sum_{j\in I_i} q_{i,j} = p$ and, for every modeling $\mathbf{A}$ with component relation ϖ and countable set of connected components $\{\mathbf{A}_k\}_{k\in\Gamma}$, it holds

$$\langle \phi, \mathbf{A}\rangle = \sum_{i=1}^{n} \prod_{j\in I_i} \sum_{k\in\Gamma} \nu_{\mathbf{A}}(A_k)^{q_{i,j}} \langle \xi_{i,j}, \mathbf{A}_k\rangle.$$

PROOF. This is a direct consequence of Lemmas 2.23 and 3.16. □

The case of sentences can be handled easily by limited counting. For a set X and an integer m, define

$$\mathrm{Big}_m(X) = \begin{cases} 1 & \text{if } |X| \geq m \\ 0 & \text{otherwise} \end{cases}$$

LEMMA 3.18. *Let $\theta \in \mathrm{FO}_0(\lambda)$.*

Then there exist formulas $\psi_1, \dots, \psi_s \in \mathrm{FO}_1^{\mathrm{local}}$ with quantifier rank at most $q(\mathrm{qrank}(\theta))$, integers $m_1, \dots, m_s \leq \mathrm{qrank}(\theta)$, and a Boolean function F such that for every λ-structure $\mathbf{A}$ with component relation ϖ and connected components $\mathbf{B}_i$ $(i \in I)$, the property $\mathbf{A} \models \theta$ is equivalent to

$$F(\mathrm{Big}_{m_1}(\{i, \mathbf{B}_i \models (\exists x)\psi_1(x)\}), \dots, \mathrm{Big}_{m_s}(\{i, \mathbf{B}_i \models (\exists x)\psi_s(x)\})) = 1.$$

PROOF. Indeed, it follows from Gaifman locality theorem 2.19 that — in presence of a component relation ϖ — every sentence θ with quantifier rank r can be written as a Boolean combination of sentences θ_k of the form

$$\exists y_1 \dots \exists y_{m_k} \Big(\bigwedge_{1 \leq i < j \leq m_k} \neg\varpi(y_i, y_j) \wedge \bigwedge_{1 \leq i \leq m_k} \psi_k(y_i) \Big)$$

where ψ_k is ϖ-local, $m_k \leq \mathrm{qrank}(\theta)$, and $\mathrm{qrank}(\psi_k) \leq q(\mathrm{qrank}(\theta))$, for some fixed function q. As $\mathbf{A} \models \theta_k$ if and only if $\mathrm{Big}_{m_k}(\{i, \mathbf{B}_i \models (\exists x)\psi_k(x)\}) = 1$, the lemma follows. $\square$

3.2.4. Convex Combinations of Modelings. In several contexts, it is clear when disjoint union of converging sequences form a converging sequence. If two graph sequences $(G_n)_{n\in\mathbb{N}}$ and $(H_n)_{n\in\mathbb{N}}$ are L-convergent or BS-convergent, it is clear that the sequence $(G_n \cup H_n)_{n\in\mathbb{N}}$ is also convergent, provided that the limit

$$\lim_{n\to\infty} |G_n|/(|G_n| + |H_n|)$$

exists. The same applies if we merge a countable set of L-convergent (resp. BS-convergent) sequences $(H_{n,i})_{n\in\mathbb{N}}$ (where $i \in \mathbb{N}$), with the obvious restriction that for each $i \in \mathbb{N}$ all but finitely many $H_{n,i}$ are empty graphs.

We shall see that the possibility to merge a countable set of converging sequences to $\mathrm{FO}^{\mathrm{local}}$-convergence will need a further assumption, namely the following equality:

$$\sum_i \lim_{n\to\infty} \frac{|G_{n,i}|}{|\bigcup_j G_{n,j}|} = 1.$$

The importance of this assumption is illustrated by the next example.

EXAMPLE 3.19. Let $N_n = 2^{2^n}$ (so that $N(n)$ is divisible by 2^i for every $1 \leq i \leq 2^n$). Consider sequences $(H_{n,i})_{n\in\mathbb{N}}$ of edgeless black and white colored graphs where $H_{n,i}$ is

- empty if $i > 2^n$,
- the edgeless graph with $(2^{-i} + 2^{-n})N_n$ white vertices and $2^{-i}N_n$ black vertices if n is odd,
- the edgeless graph with $(2^{-i} + 2^{-n})N_n$ black vertices and $2^{-i}N_n$ white vertices if n is even.

For each $i \in \mathbb{N}$, the sequence $(H_{n,i})_{n\in\mathbb{N}}$ is obviously L-convergent (and even FO-convergent) as the proportion of white vertices in $H_{n,i}$ tends to $1/2$ as $n \to \infty$. The order of $G_n = \bigcup_{i\in\mathbb{N}} H_{n,i}$ is $3N_n$ and $|H_{n,i}|/|G_n|$ tends to $\frac{2}{3}.2^{-i}$ as n goes to infinity. However the sequence $(G_n)_{n\in\mathbb{N}}$ is not L-convergent (hence not $\mathrm{FO}^{\mathrm{local}}$-convergent). Indeed, the proportion of white vertices in G_n is $2/3$ if n is odd and $1/3$ is n is even.

Hence, we are led to the following definition.

Definition 3.20 (Convex combination of Modelings). Let $\mathbf{H}_i$ be λ-modelings for $i \in I \subseteq \mathbb{N}$ and let $(\alpha_i)_{i\in I}$ be positive real numbers such that $\sum_{i\in I} \alpha_i = 1$.

Let $\mathbf{H} = \coprod_{i\in I} \mathbf{H}_i$ be the relational sample space obtained as the disjoint union of the $\mathbf{H}_i$. We endow $\mathbf{H}$ with the probability measure $\nu_{\mathbf{H}}(X) = \sum_i \alpha_i \nu_{\mathbf{H}_i}(X \cap H_i)$.

Then $\mathbf{H}$ is the *convex combination* of modelings $\mathbf{H}_i$ with *weights* α_i and we denote it by $\coprod_{i\in I}(\mathbf{H}_i, \alpha_i)$.

Lemma 3.21. *Let $\mathbf{H}_i$ be λ-modelings for $i \in I \subseteq \mathbb{N}$ and let $(\alpha_i)_{i\in I}$ be positive real numbers such that $\sum_{i\in I} \alpha_i = 1$. Let $\mathbf{H} = \coprod_{i\in I}(\mathbf{H}_i, \alpha_i)$ Then*

(1) $\mathbf{H}$ is a modeling, each H_i is measurable and $\nu_{\mathbf{H}}(H_i) = \alpha_i$ holds for every $i \in I$;
(2) if all the $\mathbf{H}_i$ are weakly uniform and either all the H_i are infinite or all the H_i are finite, I is finite, and $\alpha_i = |H_i|/\sum_{i\in I}|H_i|$, then $\mathbf{H}$ is weakly uniform.

Proof. According to Lemma 3.2, $\mathbf{H}$ is a relational sample space, in which each H_i is measurable. That $\nu_{\mathbf{H}}(H_i) = \alpha_i$ immediately follows from the definition of $\nu_{\mathbf{H}}$.

Assume that all the $\mathbf{H}_i$ are weakly uniform. If all the $\mathbf{H}_i$ are finite, I is finite, and $\alpha_i = |H_i|/\sum_{i\in I}|H_i|$, then $\mathbf{H}$ is the modeling associated to the union of the $\mathbf{H}_i$ hence it is weakly uniform. Otherwise all the H_i are infinite, hence all the $\nu_{\mathbf{H}_i}$ are atomless, $\nu_{\mathbf{H}}$ is atomless, and $\mathbf{H}$ is weakly uniform. □

Lemma 3.22. *Let $p \in \mathbb{N}$ and $\phi \in \mathrm{FO}_p^{\mathrm{local}}(\lambda)$.*

Then there exist local formulas $\xi_{i,j} \in \mathrm{FO}_{q_{i,j}}^{\mathrm{local}}$ ($1 \le i \le n$, $j \in I_i$) with $\mathrm{qrank}(xi_{i,j}) \le \mathrm{qrank}(\phi)$ such that for each i, $\sum_{j\in I_i} q_{i,j} = p$ and, for every countable set of modelings $\mathbf{A}_k$ and weights α_k ($k \in \Gamma \subseteq \mathbb{N}$ and $\sum_k \alpha_k = 1$) it holds, denoting $\mathbf{A} = \coprod_{i\in\Gamma}(\mathbf{A}_i, \alpha_i)$:

$$\langle \phi, \mathbf{A}\rangle = \sum_{i=1}^{n} \prod_{j\in I_i} \sum_{k\in\Gamma} \alpha_k^{q_{i,j}} \langle \xi_{i,j}, \mathbf{A}_k\rangle.$$

Proof. Considering, as above, the combination $\mathbf{A}^+ = \coprod_{i\in\Gamma}(\mathbf{A}_i^+, \alpha_i)$, where $\mathbf{A}_i^+$ is obtained by the basic interpretation scheme adding a full binary relation ϖ, the result is an immediate consequence of Theorem 3.17. □

For λ-modelings $\mathbf{A}$ and $\mathbf{B}$, and $p, r \in \mathbb{N}$ define

$$\|\mathbf{A} - \mathbf{B}\|_{p,r}^{\mathrm{local}} = \sup\{|\langle \phi, \mathbf{A}\rangle - \langle \phi, \mathbf{B}\rangle| :\ \phi \in \mathrm{FO}_p^{\mathrm{local}}(\lambda), \mathrm{qrank}(\phi) \le r\}.$$

The following lemma relates precisely how close are Stone pairings on two combinations of modelings, when the modelings and weights involved in the combinations define close Stone pairings.

Lemma 3.23. *Let $p, r \in \mathbb{N}$, and let $\Gamma \subseteq \mathbb{N}$. For $k \in \Gamma$, let $\mathbf{A}_k, \mathbf{B}_k$ be λ-modelings, and let α_k, β_k be non-negative weights with $\sum_k \alpha_k = \sum_k \beta_k = 1$.*

Let $\mathbf{A} = \coprod_{i\in\Gamma}(\mathbf{A}_i, \alpha_i)$ and $\mathbf{B} = \coprod_{i\in\Gamma}(\mathbf{B}_i, \beta_i)$. Then there exists a constant $c_{r,p}$ (which depends only on λ, r, and p) such that it holds

$$\begin{aligned}\|\mathbf{A}-\mathbf{B}\|_{p,r}^{\mathrm{local}} &\le c_{r,p}\big(\|\alpha-\beta\|_1 + \sum_{k\in\Gamma}\alpha_k\|\mathbf{A}_k-\mathbf{B}_k\|_{p,r}^{\mathrm{local}}\big)\\ &\le c_{r,p}\big(\|\alpha-\beta\|_1 + \sup_{i\in\Gamma}\|\mathbf{A}_i-\mathbf{B}_i\|_{p,r}^{\mathrm{local}}\big).\end{aligned}$$

Proof. Let $\phi \in \mathrm{FO}_p^{\mathrm{local}}(\lambda)$ with $\mathrm{qrank}(\phi) \le r$. According to Lemma 3.22 there exist local formulas $\xi_{i,j} \in \mathrm{FO}_{q_{i,j}}^{\mathrm{local}}(\lambda)$ $(1 \le i \le n,\ j \in I_i)$ with $\mathrm{qrank}(xi_{i,j}) \le r$ such that for each i, $\sum_{j\in I_i} q_{i,j} = p$ and, for every countable set of modelings $\mathbf{C}_k$ and weights γ_k ($k \in \Gamma$ and $\sum_k \gamma_k = 1$) it holds, denoting $\mathbf{C} = \coprod_{i\in\Gamma}(\mathbf{C}_i, \gamma_i)$:

$$\langle\phi, \mathbf{C}\rangle = \sum_{i=1}^n \prod_{j\in I_i}\langle\xi_{i,j}, \mathbf{C}\rangle, \quad \text{with } \langle\xi_{i,j}, \mathbf{C}\rangle = \sum_{k\in\Gamma}\gamma_k^{q_{i,j}}\langle\xi_{i,j}, \mathbf{C}_k\rangle.$$

As there are only finitely many non-equivalent formulas in $\mathrm{FO}_p^{\mathrm{local}}(\lambda)$ with quantifier rank at most r, there is a constant $N_{r,p}$ such that $n \le N_{r,p}$.

We have

$$|\langle\phi, \mathbf{A}\rangle - \langle\phi, \mathbf{B}\rangle| \le \sum_{i=1}^n\left|\prod_{j\in I_i}\langle\xi_{i,j}, \mathbf{A}\rangle - \prod_{j\in I_i}\langle\xi_{i,j}, \mathbf{B}\rangle\right|$$

Note that if $a_i, b_i \in [0, 1]$ then we get easily

$$\begin{aligned}\left|\prod_{i=1}^k a_i - \prod_{i=1}^k b_i\right| &= \left|(a_1-b_1)\prod_{i=2}^k a_i + b_1\big(\prod_{i=2}^k a_i - \prod_{i=2}^k b_i\big)\right|\\ &\le |a_1-b_1| + \left|\prod_{i=2}^k a_i - \prod_{i=2}^k b_i\right|\\ &\le \sum_{i=1}^k |a_i-b_i|\end{aligned}$$

Hence, as for every $1 \le i \le n$ and every $j \in I$ it holds $0 \le \langle\xi_{i,j}, \mathbf{A}\rangle \le 1$ and $0 \le \langle\xi_{i,j}, \mathbf{B}\rangle \le 1$, we have

$$\begin{aligned}|\langle\phi, \mathbf{A}\rangle - \langle\phi, \mathbf{B}\rangle| &\le \sum_{i=1}^n\sum_{j\in I_i}\Big|\langle\xi_{i,j}, \mathbf{A}\rangle - \langle\xi_{i,j}, \mathbf{B}\rangle\Big|\\ &\le \sum_{i=1}^n\sum_{j\in I_i}\Big|\sum_{k\in\Gamma}\alpha_k^{q_{i,j}}\langle\xi_{i,j}, \mathbf{A}_k\rangle - \sum_{k\in\Gamma}\beta_k^{q_{i,j}}\langle\xi_{i,j}, \mathbf{B}_k\rangle\Big|\\ &\le \sum_{i=1}^n\sum_{j\in I_i}\sum_{k\in\Gamma}\big|\alpha_k^{q_{i,j}}\langle\xi_{i,j}, \mathbf{A}_k\rangle - \beta_k^{q_{i,j}}\langle\xi_{i,j}, \mathbf{B}_k\rangle\big|\end{aligned}$$

Thus, as $q_{i,j} \geq 1$ and as Stone pairings $\langle \cdot, \cdot \rangle$ have value in $[0,1]$, it holds (denoting $c_{r,p} = pN_{r,p}$):

$$\|\mathbf{A}-\mathbf{B}\|_{p,r}^{\mathrm{local}} \leq c_{r,p}\left(\sum_{k\in\Gamma} |\alpha_k - \beta_k| + \sum_{k\in\Gamma} \alpha_k \|\mathbf{A}_k - \mathbf{B}_k\|_{p,r}^{\mathrm{local}}\right)$$

□

LEMMA 3.24. *Let $p, r \in \mathbb{N}$, let $\mathbf{A}, \mathbf{B}$ be λ-modeling, with connected components $\mathbf{A}_k, k \in \Gamma_{\mathbf{A}}$ and $\mathbf{B}_k, k \in \Gamma_{\mathbf{B}}$ (where $\Gamma_{\mathbf{A}}$ and $\Gamma_{\mathbf{B}}$ can be infinite non-countable).*

Then it holds

$$\|\mathbf{A}-\mathbf{B}\|_{p,r}^{\mathrm{local}} < c_{r,p}\big(\sup_{k\in\Gamma_{\mathbf{A}}} \nu_{\mathbf{A}}(\mathbf{A}_k) + \sup_{k\in\Gamma_{\mathbf{B}}} \nu_{\mathbf{B}}(\mathbf{B}_k) + \|\mathbf{A}-\mathbf{B}\|_{1,r}^{\mathrm{local}}\big).$$

PROOF. Let $\phi \in \mathrm{FO}_p^{\mathrm{local}}(\lambda)$ with $\mathrm{qrank}(\phi) \leq r$. It holds

$$\langle \phi, \mathbf{A}\rangle = \sum_{i=1}^{n} \prod_{j\in I_i} \langle \xi_{i,j}, \mathbf{A}\rangle.$$

It is clear that if $\zeta_{i,j}$ is component-local and $q_{i,j} > 1$ then

$$\langle \xi_{i,j}, \mathbf{A}\rangle < \sup_{k\in\Gamma_{\mathbf{A}}} \nu_{\mathbf{A}}(\mathbf{A}_k).$$

Let X be the set of the integers $1 \leq i \leq n$ such that there is $j \in I_i$ such that $q_{i,j} > 1$, and let Y be the complement of X in $\{1, \ldots, n\}$. Then

$$\Big|\langle \phi, \mathbf{A}\rangle - \sum_{i\in Y} \prod_{j\in I_i} \langle \xi_{i,j}, \mathbf{A}\rangle\Big| < c_{r,p} \sup_{k\in\Gamma_{\mathbf{A}}} \nu_{\mathbf{A}}(\mathbf{A}_k).$$

Similarly, it holds

$$\Big|\langle \phi, \mathbf{B}\rangle - \sum_{i\in Y} \prod_{j\in I_i} \langle \xi_{i,j}, \mathbf{B}\rangle\Big| < c_{r,p} \sup_{k\in\Gamma_{\mathbf{B}}} \nu_{\mathbf{B}}(\mathbf{B}_k).$$

Thus the statement follows from

$$\begin{aligned}\left|\sum_{i\in Y} \prod_{j\in I_i} \langle \xi_{i,j}, \mathbf{A}\rangle - \sum_{i\in Y} \prod_{j\in I_i} \langle \xi_{i,j}, \mathbf{B}\rangle\right| &\leq \sum_{i\in Y} \left|\prod_{j\in I_i} \langle \xi_{i,j}, \mathbf{A}\rangle - \prod_{j\in I_i} \langle \xi_{i,j}, \mathbf{B}\rangle\right| \\ &\leq \sum_{i\in Y} \sum_{j\in I_i} \big|\langle \xi_{i,j}, \mathbf{A}\rangle - \langle \xi_{i,j}, \mathbf{B}\rangle\big| \\ &\leq c_{r,p} \|\mathbf{A}-\mathbf{B}\|_{1,r}^{\mathrm{local}}\end{aligned}$$

□

THEOREM 3.25. *Let $p \in \mathbb{N}$, let $I \subseteq \mathbb{N}$ and, for each $i \in I$ let $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ be an $\mathrm{FO}_p^{\mathrm{local}}(\lambda)$-convergent sequence of λ-modelings and let $(a_{i,n})_{n\in\mathbb{N}}$ be a convergent sequence of non-negative real numbers, such that $\sum_{i\in I} a_{i,n} = 1$ holds for every $n \in \mathbb{N}$, and such that $\sum_{i\in I} \lim_{n\to\infty} a_{i,n} = 1$.*

Then the sequence of convex combinations $\coprod_{i\in I}(\mathbf{A}_{i,n}, a_{i,n})$ is $\mathrm{FO}_p^{\mathrm{local}}(\lambda)$-convergent.

Proof. If I is finite, then the result follows from Lemma 3.22. Hence we can assume $I = \mathbb{N}$.

Let $\phi \in \mathrm{FO}_p^{\mathrm{local}}$, let $q \in \mathbb{N}$, and let $\epsilon > 0$ be a positive real. Assume that for each $i \in \mathbb{N}$ the sequence $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ is $\mathrm{FO}_p^{\mathrm{local}}$-convergent and that $(a_{i,n})_{n\in\mathbb{N}}$ is a convergent sequence of non-negative real numbers, such that $\sum_i a_{i,n} = 1$ holds for every $n \in \mathbb{N}$. Let $\alpha_i = \lim_{n\to\infty} a_{i,n}$, let $d_i = \lim_{n\to\infty}\langle \phi, \mathbf{A}_{i,n}\rangle$, and let C be such that $\sum_{i=1}^C \alpha_i > 1 - \epsilon/4$. There exists N such that for every $n \geq N$ and every $i \leq C$ it holds $|a_{n,i} - \alpha_i| < \epsilon/4C$ and $|a_{i,n}^q \langle \phi, \mathbf{A}_{i,n}\rangle - \alpha_i^q d_i| < \epsilon/2C$. Thus $\left|\sum_{i=1}^C a_{i,n}^q \langle \phi, \mathbf{A}_{i,n}\rangle - \sum_{i=1}^C \alpha_i^q d_i\right| < \epsilon/2$ and $\sum_{i>C+1} a_{i,n} < \epsilon/2$. It follows that for any $n \geq N$ it holds

$$\left|\sum_{i>C+1} a_{i,n}^q \langle \phi, \mathbf{A}_{i,n}\rangle - \sum_{i>C+1} \alpha_i^q d_i\right| \leq \max\left(\sum_{i>C+1} a_{i,n}^q, \sum_{i>C+1} \alpha_i^q d_i\right) < \epsilon/2$$

hence $|\sum_i a_{i,n}^q \langle \phi, \mathbf{A}_{i,n}\rangle - \sum_i \alpha_i^q d_i| < \epsilon$.

For every $\psi \in \mathrm{FO}_p^{\mathrm{local}}$, the expression appearing in Lemma 3.22 for the expansion of $\langle \phi, \coprod_i (\mathbf{A}_{i,n}, a_{i,n})\rangle$ is a finite combination of terms of the form $\sum_i a_{i,n}^{q_{i,n}} \langle \phi, \mathbf{A}_{i,n}\rangle$, where $q_{i,n} \in \mathbb{N}$ and $\phi \in \mathrm{FO}_p^{\mathrm{local}}$. It follows that the value $\langle \phi, \coprod_i (\mathbf{A}_{i,n}, a_{i,n})\rangle$ converges as n grows to infinity. Hence $(\coprod_i (\mathbf{A}_{i,n}, a_{i,n}))_{n\in\mathbb{N}}$ is $\mathrm{FO}_p^{\mathrm{local}}$-convergent. □

Corollary 3.4. *Let $p \geq 1$ and let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be a sequence of finite λ-structures. Assume $\mathbf{A}_n$ be the disjoint union of $\mathbf{B}_{n,i}$ ($i \in \mathbb{N}$) where all but a finite number of $\mathbf{B}_{n,i}$ are empty. Let $a_{n,i} = |B_{n,i}|/|A_n|$. Assume further that:*

- *for each $i \in \mathbb{N}$, the limit $\alpha_i = \lim_{n\to\infty} a_{n,i}$ exists,*
- *for each $i \in \mathbb{N}$ such that $\alpha_i \neq 0$, the sequence $(\mathbf{B}_{n,i})_{n\in\mathbb{N}}$ is $\mathrm{FO}_p^{\mathrm{local}}$-convergent,*
- *it holds*
$$\sum_{i\geq 1} \alpha_i = 1.$$

Then, the sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}_p^{\mathrm{local}}$-convergent.

Moreover, if $\mathbf{L}_i$ is a modeling $\mathrm{FO}_p^{\mathrm{local}}$-limit of $(\mathbf{B}_{n,i})_{n\in\mathbb{N}}$ when $\alpha_i \neq 0$ then $\coprod_i (\mathbf{L}_i, \alpha_i)$ is a modeling $\mathrm{FO}_p^{\mathrm{local}}$-limit of $(\mathbf{A}_n)_{n\in\mathbb{N}}$.

Proof. This follows from Theorem 3.25, as $\mathbf{A}_n = \coprod_i (\mathbf{B}_{n,i}, a_{n,i})$. □

Definition 3.26. A family of sequence $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ ($i \in I$) of λ-structures is *uniformly elementarily convergent* if, for every formula $\phi \in \mathrm{FO}_1(\lambda)$ there is an integer N such that it holds

$$\forall i \in I,\ \forall n' \geq n \geq N, \quad (\mathbf{A}_{i,n} \models (\exists x)\phi(x)) \Longrightarrow (\mathbf{A}_{i,n'} \models (\exists x)\phi(x)).$$

First notice that if a family $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ ($i \in I$) of sequences is uniformly elementarily convergent, then each sequence $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ is elementarily convergent

Lemma 3.27. *Let $I \subseteq N$, and let $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ ($i \in I$) be sequences forming a uniformly elementarily convergent family.*

Then $(\bigcup_{i\in I} \mathbf{A}_{i,n})_{n\in\mathbb{N}}$ is elementarily convergent.

Moreover, if $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ *is elementarily convergent to* $\widehat{\mathbf{A}}_i$ *then* $(\bigcup_{i\in I}\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ *is elementarily convergent to* $\bigcup_{i\in I}\widehat{\mathbf{A}}_i$.

Proof. Let λ^+ be the signature λ augmented by a binary relational symbol ϖ. Let I_1 be the basic interpretation scheme of λ^+-structures in λ-structures defining $\varpi(x,y)$ for every x,y. Let $\mathbf{A}_{i,n}^+ = I_1(\mathbf{A}_{i,n})$. According to Lemma 3.18, for every sentence $\theta\in\mathrm{FO}_0(\lambda)$ there exist formulas $\psi_1,\dots,\psi_s\in\mathrm{FO}_1^{\mathrm{local}}$, an integer m, and a Boolean function F such that the property $\bigcup_{i\in I}\mathbf{A}_{i,n}^+\models\theta$ is equivalent to

$$F(\mathrm{Big}_{m_1}(\{i,\mathbf{A}_{i,n}\models(\exists x)\psi_1(x)\}),\dots,\mathrm{Big}_{m_s}(\{i,\mathbf{A}_{i,n}\models(\exists x)\psi_s(x)\}))=1.$$

According to the definition of a uniformly elementarily convergent family there is an integer N such that, for every $1\le j\le s$, the value $\mathrm{Big}_{m_j}(\{i,\mathbf{A}_{i,n}\models(\exists x)\psi_j(x)\})$ is a function of n, which is non-decreasing for $n\ge N$. It follows that this function admits a limit for every $1\le j\le s$ hence the exists an integer N' such that either $\bigcup_{i\in I}\mathbf{A}_{i,n}^+\models\theta$ holds for every $n\ge N'$ or it holds for no $n\ge N'$. It follows that $(\bigcup_{i\in I}\mathbf{A}_{i,n}^+)_{n\in\mathbb{N}}$ is elementarily convergent. Thus (by means of the basic interpretation scheme deleting ϖ) $(\bigcup_{i\in I}\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ is elementarily convergent

If I is finite, it is easily checked that if $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ is elementarily convergent to $\widehat{\mathbf{A}}_i$ then $(\bigcup_{i\in I}\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ is elementarily convergent to $\bigcup_{i\in I}\widehat{\mathbf{A}}_i$.

Otherwise, we can assume $I=\mathbb{N}$. Following the same lines, it is easily checked that $(\bigcup_{i=1}^n\widetilde{\mathbf{A}}_i)_{n\in\mathbb{N}}$ converges elementarily to $(\bigcup_{i\in\mathbb{N}}\widetilde{\mathbf{A}}_i)_{n\in\mathbb{N}}$. For $i,n\in\mathbb{N}$, let $\mathbf{B}_{i,2n}=\mathbf{A}_{i,n}$ and $\mathbf{B}_{i,2n+1}=\widetilde{\mathbf{A}_i}$. As, for each $i\in\mathbb{N}$, $\widetilde{\mathbf{A}}_i$ is an elementary limit of $(\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ it is easily checked that the family of the sequences $(\mathbf{B}_{i,n})_{n\in\mathbb{N}}$ is uniformly elementarily convergent. It follows that $(\bigcup_{i\in\mathbb{N}}\mathbf{B}_{i,n})_{n\in\mathbb{N}}$ is elementarily convergent thus the elementary limit of $(\bigcup_{i\in I}\mathbf{A}_{i,n})_{n\in\mathbb{N}}$ and $(\bigcup_{i=1}^n\widetilde{\mathbf{A}}_i)_{n\in\mathbb{N}}$ are the same, that is $\bigcup_{i\in I}\widetilde{\mathbf{A}}_i$. □

¿From Corollary 3.4 and Lemma 3.27 then follows the next general result.

Corollary 3.5. *Let* $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *be a sequence of finite* λ*-structures.*

Assume $\mathbf{A}_n$ *be the disjoint union of* $\mathbf{B}_{n,i}$ $(i\in\mathbb{N})$ *where all but a finite number of* $\mathbf{B}_{n,i}$ *are empty. Let* $a_{n,i}=|B_{n,i}|/|A_n|$. *Assume that:*

- *for each* $i\in\mathbb{N}$, *the limit* $\alpha_i=\lim_{n\to\infty}a_{n,i}$ *exists and it holds*
$$\sum_{i\ge1}\alpha_i=1,$$
- *for each* $i\in\mathbb{N}$ *such that* $\alpha_i\neq0$, *the sequence* $(\mathbf{B}_{n,i})_{n\in\mathbb{N}}$ *is* $\mathrm{FO}^{\mathrm{local}}$*-convergent,*
- *the family* $\{(\mathbf{B}_{n,i})_{n\in\mathbb{N}}\ (i\in\mathbb{N})\}$ *is uniformly elementarily convergent.*

Then, the sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *is* FO*-convergent.*

Moreover, if $\mathbf{L}_i$ *is a modeling* FO*-limit of* $(\mathbf{B}_{n,i})_{n\in\mathbb{N}}$ *when* $\alpha_i\neq0$ *and an elementary limit of* $(\mathbf{B}_{n,i})_{n\in\mathbb{N}}$ *when* $\alpha_i=0$ *then* $\coprod_i(\mathbf{L}_i,\alpha_i)$ *is a modeling* FO*-limit of* $(\mathbf{A}_n)_{n\in\mathbb{N}}$.

3.2.5. Random-free graphons and Modelings. A graphon is *random-free* if it is $\{0,1\}$-valued almost everywhere. Moreover, if two graphons represent the same L-limit of finite graphs, then either they are both random-free or none of them are (see for instance [44]). Several properties of random-free graph limits have been studied.

For example, a graph limit Γ is random-free if and only if the random graph $G(n,\Gamma)$ of order n sampled from Γ has entropy $o(n^2)$ [4, 44] (see also [39]).

A sequence of graphs $(G_n)_{n\in\mathbb{N}}$ is L-convergent to a random-free graphon if and only if the sequence $(G_n)_{n\in\mathbb{N}}$ is convergent for the stronger metric δ_1 [73], where the distance $\delta_1(G,H)$ of graphs G and H with respective vertex sets $\{x_1,\dots,x_m\}$ and $\{y_1,\dots,y_m\}$ is the minimum over all non-negative $m\times n$ matrices $A=(\alpha_{i,j}$ with row sums $1/m$ and column sums $1/n$ of $\sum_{(i,j,g,h)\in\Delta}\alpha_{i,g}\alpha_{j,h}$, where Δ is the set of quadruples (i,j,g,h) such that either $\{x_i,x_j\}\in E(G)$ or $\{y_g,y_h\}\in E(H)$ (but not both).

Lovász and Szegedy [55] defined a graph property (or equivalently a class of graphs) $\mathcal{C}$ to be *random-free* if every L-limit of graphs in $\mathcal{C}$ is random-free. They prove the following:

THEOREM 3.28 (Lovász and Szegedy [55]). *A hereditary class $\mathcal{C}$ is random-free if and only if there exists a bipartite graph F with bipartition (V_1,V_2) such that no graph obtained from F by adding edges within V_1 and V_2 is in $\mathcal{C}$.*

¿From this result, one deduce for instance that the class of m-partite cographs is random-free [17], generalizing the particular cases of threshold graphs [20] and (more general) cographs [43].

Recall that the *Vapnik–Chervonenkis dimension* (or simply *VC-dimension*) $\mathrm{VC}(G)$ *of a graph* G is the maximum integer k such that there exists in G disjoint vertices u_i $(1\le i\le k)$ and v_I $(\emptyset\subseteq I\subseteq\{1,\dots,k\})$ such that u_i is adjacent to v_I exactly if $i\in I$. We now rephrase Lovász and Szegedy Theorem 3.28 in terms of VC-dimension.

THEOREM 3.29. *A hereditary class $\mathcal{C}$ is random-free if and only if* $\mathrm{VC}(\mathcal{C})<\infty$, *where*

$$\mathrm{VC}(\mathcal{C})=\sup_{G\in\mathcal{C}}\mathrm{VC}(G).$$

PROOF. Let S_k be the bipartite graph with vertices u_i $(1\le i\le k)$ and v_I $(\emptyset\subseteq I\subseteq\{1,\dots,k\})$ such that u_i is adjacent to v_I exactly if $i\in I$.

If $\mathrm{VC}(\mathcal{C})<k$ then no graph obtained from S_k by adding edges within the u_i's and the v_I's is in $\mathcal{C}$ hence, according to Theorem 3.28, the class $\mathcal{C}$ is random-free.

Conversely, if the class $\mathcal{C}$ is random-free there exists, according to Theorem 3.28, a bipartite graph F with bipartition (V_1,V_2) (with $|V_1|\le|V_2|$) such that no graph obtained from F by adding edges within V_1 and V_2 is in $\mathcal{C}$. It is easily checked that F is an induced subgraph of $S_{|V_1|+\log_2|V_2|}$ so $\mathrm{VC}(\mathcal{C})<\frac{|F|}{2}+\log_2|F|$. □

The VC-dimension of classes of graphs can also be related to the nowhere dense/somewhere dense dichotomy. Recall that a class $\mathcal{C}$ is *somewhere dense* if there exists an integer p such that for every integer n the p-subdivision of K_n is a subgraph of a graph in $\mathcal{C}$, and that the class $\mathcal{C}$ is *nowhere dense*, otherwise [65, 69, 71]. This dichotomy can also be characterized in quite a number of different ways, see [71]. Based on Laskowski [49], another characterization has been proved, which relates this dichotomy to VC-dimension:

THEOREM 3.30 (Adler and Adler [2]). *For a monotone class of graphs $\mathcal{C}$, the following are equivalent:*

(1) For every interpretation scheme I *of graphs in graphs, the class* $\mathsf{I}(\mathcal{C})$ *has bounded VC-dimension;*
(2) For every basic interpretation scheme I *of graphs in graphs with exponent* 1*, the class* $\mathsf{I}(\mathcal{C})$ *has bounded VC-dimension;*
(3) The class $\mathcal{C}$ *is nowhere-dense.*

¿From Theorem 3.28 and 3.30 we deduce the following:

THEOREM 3.31. *Let* $\mathcal{C}$ *be a monotone class of graphs. Then the following are equivalent:*

(1) For every interpretation scheme I *of graphs in graphs, the class* $\mathsf{I}(\mathcal{C})$ *is random-free;*
(2) For every basic interpretation scheme I *of graphs in graphs with exponent* 1*, the class* $\mathsf{I}(\mathcal{C})$ *is random-free;*
(3) The class $\mathcal{C}$ *is nowhere-dense.*

PROOF. Obviously, condition (1) implies condition (2). Assume that (2) and assume for contradiction that (3) does not hold. Then, as $\mathcal{C}$ is monotone and somewhere dense, there is an integer $p \geq 1$ such that for every graph n, the p-subdivision $\mathrm{Sub}_p(K_n)$ of the complete graph K_n is in $\mathcal{C}$. To every finite graph G we associate a graph G' by considering an arbitrary orientation of G and then building G' as shown on the figure bellow.

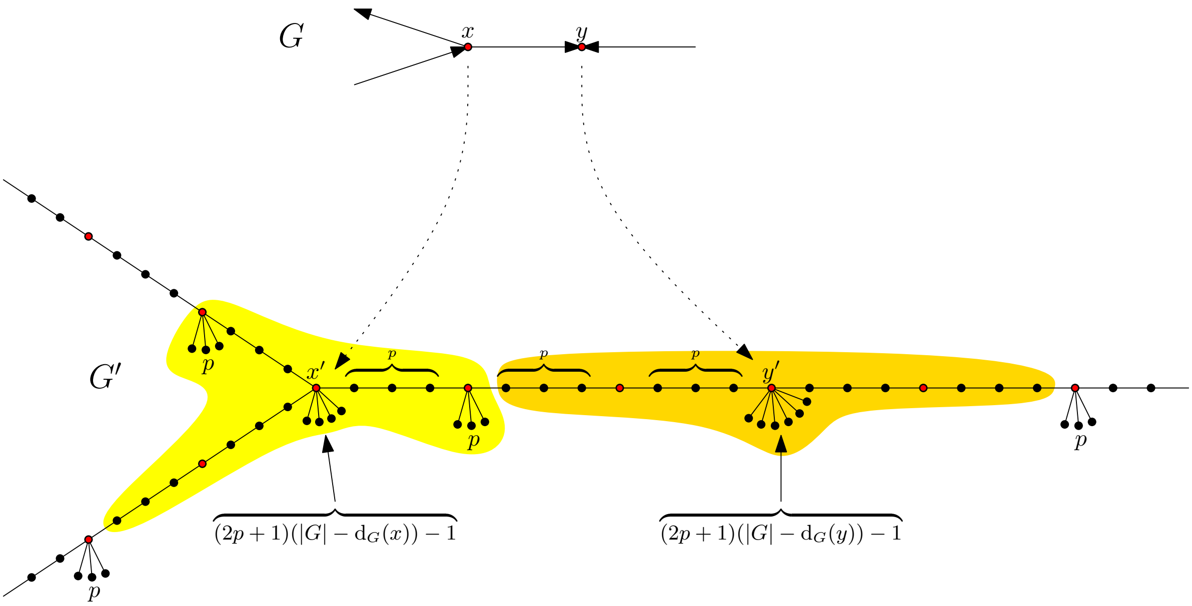

Note that $G' \in \mathcal{C}$ as it is obviously a subgraph of the p-subdivision of the complete graph of order $(2p+1)|G|^2$. It is easily checked that there is a basic interpretation scheme I_p of graphs in graphs with exponent 1 (which definitions only depends on p) such that $\mathsf{I}_p(G') = G[(2p+1)|G|]$, where $G[(2p+1)|G|]$ denotes the graph obtained from G by blowing each vertex to an independent set of size $(2p+1)|G|$.

Let $(G_i)_{i\in\mathbb{N}}$ be a sequence of graph that is L-convergent to a non random-free graphon W. As $t(F,G) = \frac{\hom(F,G)}{|G|^{|F|}}$ is invariant by uniform blow-up of the vertices of G, for every finite graph F it holds

$$t(F, \mathsf{I}_p(G_i')) = t(F, G_i[(2p+1)|G_i|]) = t(F, G_i).$$

Hence $(\mathsf{I}_p(G'_i))_{i\in\mathbb{N}}$ is L-convergent to W. Then the condition 2 contradicts the hypothesis that W is not random-free. It follows (by contradiction) that (2) implies (3).

Assume condition (3) holds, and let I be an interpretation scheme of graphs in graphs. Then according to Theorem 3.30 the class $\mathsf{I}(\mathcal{C})$ has bounded VC-dimension, hence the hereditary closure of $\mathsf{I}(\mathcal{C})$ has bounded VC-dimension thus is random-free, whence $\mathsf{I}(\mathcal{C})$ is random-free. □

We derive the following corollary concerning existence of modeling FO-limits, which completes Corollary 3.2.

Theorem 3.32. *Let $\mathcal{C}$ be a monotone class.*
If every FO*-convergent sequence of graphs in $\mathcal{C}$ has a modeling* FO*-limit then $\mathcal{C}$ is nowhere dense.*

Proof. Let I be a basic interpretation scheme of graphs in graphs, and let $(G_i)_{i\in\mathbb{N}}$ be a sequence of graphs in $\mathcal{C}$ such that $|G_i|$ is unbounded, and the sequence $(\mathsf{I}(G_i))_{i\in\mathbb{N}}$ is L-convergent.

By compactness, the sequence $(G_i)_{i\in\mathbb{N}}$ has a subsequence $(G_{n_i})_{i\in\mathbb{N}}$ that is FO-convergent. Hence, by hypothesis, $(G_{n_i})_{i\in\mathbb{N}}$ has a modeling FO-limit $\mathbf{L}$. According to Proposition 3.1, the sequence $(\mathsf{I}(G_{n_i}))_{i\in\mathbb{N}}$ has modeling FO-limit $\mathsf{I}(\mathbf{L})$. By Lemma 3.9, $\mathbf{L}$ defines a random-free graphon W that is the L-limit of $(\mathsf{I}(G_{n_i}))_{i\in\mathbb{N}}$. Of course, the L-limit of an L-convergent sequence $(G_i)_{i\in\mathbb{N}}$ with $|G_i|$ bounded is also random-free. Hence the class $\mathsf{I}(\mathcal{C})$ is random-free. As this conclusion holds for every basic interpretation scheme I we deduce from Theorem 3.31 that $\mathcal{C}$ is nowhere dense. □

Actually, we conjecture that the converse is true (see Conjecture 1.1).

3.2.6. Modelings FO-limits for Graphs of Bounded Degrees. Nice limit objects are known for sequence of bounded degree connected graphs, both for BS-convergence (graphing) and for FO_0-convergence (countable graphs). It is natural to ask whether a nice limit object could exist for full FO-convergence. We shall now give a positive answer to this question. First we take time to comment on the connectivity assumption. A first impression is that FO-convergence of disconnected graphs could be considered component-wise. This is far from being true in general. The contrast between the behaviour of graphs with a first-order definable component relation (like graphs with bounded diameter components) and of graphs with bounded degree is exemplified by the following example.

Example 3.33. Consider a BS-convergent sequence $(G_n)_{n\in\mathbb{N}}$ of planar graphs with bounded degrees such that the limit distribution has an infinite support. Note that $\lim_{n\to\infty}|G_n| = \infty$. Then, as planar graphs with bounded degrees form a hyperfinite class of graphs there exists, for every graph G_n and every $\epsilon > 0$ a subgraph $S(G_n, \epsilon)$ of G_n obtained by deleting at most $\epsilon|G_n|$ of edges, such that the connected components of $S(G_n, \epsilon)$ have order at most $f(\epsilon)$. By considering a subsequence $G_{s(n)}$ we can assume $\lim_{n\to\infty}|G_{s(n)}|/f(1/n) = \infty$. Then note that the sequences $(G_s(n))_{n\in\mathbb{N}}$ and $(S(G_{s(n)}, 1/n))_{n\in\mathbb{N}}$ have the same BS-limit. By merging these sequences, we conclude that there exists an $\mathrm{FO}^{\mathrm{local}}$ convergent sequence of graphs with bounded degrees (H_n) such that H_n is connected if n is even and such that the number of connected components of H_n for n odd tends to infinity.

EXAMPLE 3.34. Using Fig. 2, cConsider four sequences $(A_n)_{n\in\mathbb{N}}$, $(B_n)_{n\in\mathbb{N}}$, $(C_n)_{n\in\mathbb{N}}$,$(D_n)_{n\in\mathbb{N}}$ of FO-converging sequences where $|A_n| = |B_n| = |C_n| = |D_n|$ grows to infinity, and where these sequences have distinct limits.

Consider a sequence $(G_n)_{n\in\mathbb{N}}$ defined as follows: for each n, G_n has two connected components denoted by $H_{n,1}$ and $H_{n,2}$ obtained by joining A_n, C_n and B_n, D_n by a path of length n (for n odd), and by joining A_n, D_n and B_n, C_n by a path of length n (for n even). Then $(G_n)_{n\in\mathbb{N}}$ is FO-convergent. However, there is no choice of a mapping $f : \mathbb{N} \to \{1, 2\}$ such that $(H_{n,f(n)})$ is FO-convergent (or even BS-convergent).

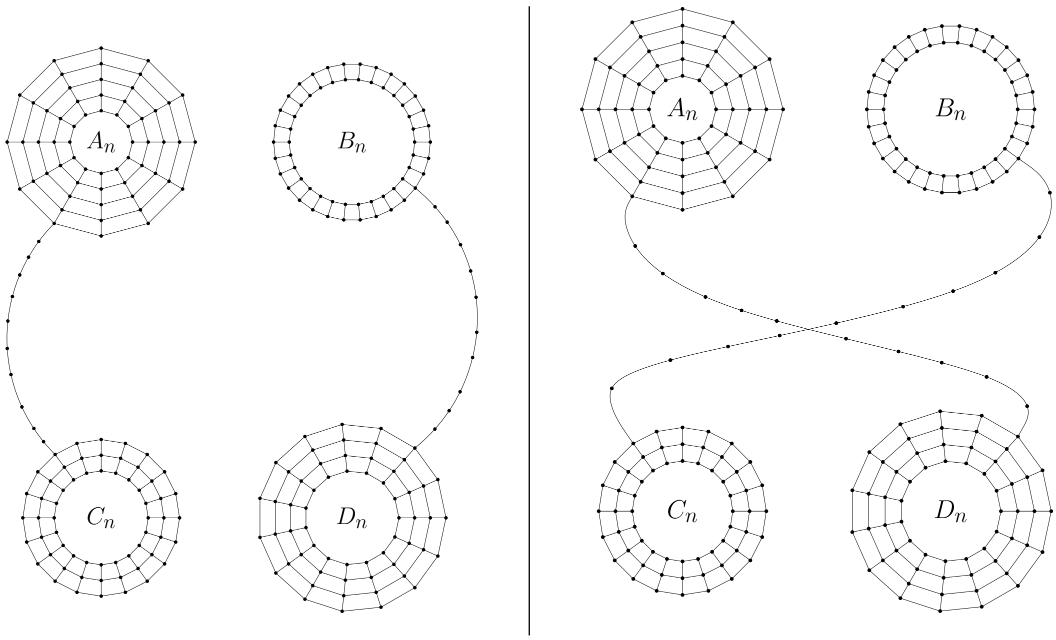

FIGURE 2. An FO-converging sequence with no component selection

This situation is indeed related to the fact that the diameter of the graph G_n in the sequence tend to infinity as n grows and that the belonging to a same connected component cannot be defined by a first-order formula. This situation is standard when one consider BS-limits of connected graphs with bounded degrees: it is easily checked that, as a limit of connected graphs, a graphing may have uncountably many connected components.

REMARK 3.35. In the spirit of the construction shown Fig. 2, we can prove that the set of measure μ which are BS-limits of connected graphs with maximum degree $d \geq 2$ and order going to infinity is convex: Assume $(G_n)_{n\in\mathbb{N}}$ and $(H_n)_{n\in\mathbb{N}}$ are convergent sequences with limits μ_1 and μ_2, and let $0 < \alpha < 1$. We construct graph M_n as follows: let $c_n = \min(|G_n|, |H_n|)$. We consider $\alpha|H_n|$ copies of G_n and $(1-\alpha)|G_n|$ copies of H_n linked by paths of length $\lfloor \log c_n \rfloor$ (see Fig. 3). It is easily checked that the statistics of the neighborhoods of M_n tend to $\alpha\mu_1 + (1-\alpha)\mu_2$.

Let V be a standard Borel space with a measure μ. Suppose that $T_1, T_2, \dots, T_k$ are measure preserving Borel involutions of X. Then the system

$$\mathbf{G} = (V, T_1, T_2, \dots, T_k, \mu)$$

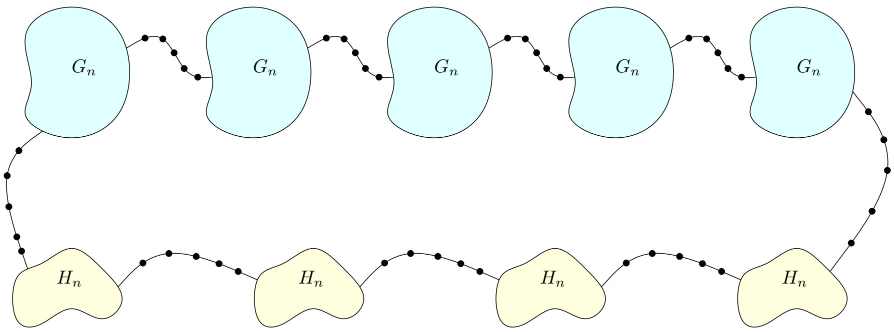

FIGURE 3. Construction of the graph M_n

is called a *measurable graphing* (or simply a *graphing*) [1]. A graphing $\mathbf{G}$ determines an equivalence relation on the points of V. Simply, $x \sim_{\mathbf{G}} y$ if there exists a sequence of points $(x_1, x_2, \ldots, x_m)$ of X such that

- $x_1 = x, x_m = y$
- $x_{i+1} = T_j(x_i)$ for some $1 \leq j \leq k$.

Thus there exist natural a simple graph structure on the equivalence classes, the *leafgraph.* Here x is adjacent to y, if $x \neq y$ and $T_j(x) = y$ for some $1 \leq j \leq k$. Now if If V is a compact metric space with a Borel measure μ and $T_1, T_2, \ldots, T_k$ are continuous measure preserving involutions of V, then $\mathbf{G} = (V, T_1, T_2, \ldots, T_k, \mu)$ is a *topological graphing.* It is a consequence of [8] and [34] that every local weak limit of finite connected graphs with maximum degree at most D can be represented as a measurable graphing. Elek [28] further proved the representation can be required to be a topological graphing.

A graphing defines an edge coloration, where $\{x, y\}$ is colored by the set of the indexes i such that $y = T_i(x)$. For an integer r, a graphing $\mathbf{G} = (V, T_1, \ldots, T_k, \mu)$ and a finite rooted edge colored graph (F, o) we define the set

$$D_r(\mathbf{G}, (F, o)) = \{x \in \mathbf{G}, B_r(\mathbf{G}, x) \simeq (F, o)\}.$$

It is easily checked that $D_r(\mathbf{G}, (F, o))$ is measurable.

Considering k-edge colored graphing allows to describe a vertex x in a distance-r neighborhood of a given vertex v by the sequence of the colors of the edges of a path linking v to x. Taking, among the minimal length sequences, the one which is lexicographically minimum, it is immediate that for every vertex v and every integer r there is a injection $\iota_{v,r}$ from $B_r(\mathbf{G}, v)$ to the set of the sequences of length at most r with values in $[k]$. Moreover, if $B_r(\mathbf{G}, v)$ and $B_r(\mathbf{G}, v')$ are isomorphic as edge-colored rooted graphs, then there exists a unique isomorphism $f : B_r(\mathbf{G}, v) \to B_r(\mathbf{G}, v')$ and this isomorphism as the property that for every $x \in B_r(\mathbf{G}, v)$ it holds $\iota_{v,r'}(f(x)) = \iota_{v,r}(x)$.

LEMMA 3.36. *Every graphing is a modeling.*

PROOF. Let $\mathbf{G} = (V, T_1, \ldots, T_d, \mu)$ be a graphing. We color the edges of G according to the the involutions involved.

For $r \in \mathbb{N}$, we denote by $\mathcal{F}_r$ the finite set of all the colored rooted graphs that arise as $B_r(\mathbf{G}, v)$ for some $v \in V$. To every vertex $v \in V$ and integer $r \in \mathbb{N}$ we associate $t_r(v)$, which is the isomorphism type of the edge colored ball $B_r(\mathbf{G}, v)$.

According to Gaifman's locality theorem, in order to prove that $\mathbf{G}$ is a modeling, it is sufficient to prove that for each $\phi \in FO_p^{\mathrm{local}}$, the set

$$X = \{(v_1, \ldots, v_p) \in V^q : \quad \mathbf{G} \models \phi(v_1, \ldots, v_p)\}$$

is measurable (with respect to the product σ-algebra of V^p).

Let $L \in \mathbb{N}$ be such that ϕ is L-local. For every $\mathbf{v} = (v_1, \ldots, v_p) \in X$ we define the graph $\Gamma(\mathbf{v})$ with vertex set $\{v_1, \ldots, v_p\}$ such that two vertices of $\Gamma(\mathbf{v})$ are adjacent if their distance in $\mathbf{G}$ is at most L. We define a partition $\mathcal{P}(\mathbf{v})$ of $[p]$ as follows: i and j are in a same part if v_i and v_j belong to a same connected component of $\Gamma(v)$. To each part $P \in \mathcal{P}(\mathbf{v})$, we associate the tuple formed by $T_P = t_{(|P|-1)L}(v_{\min P})$ and, for each $i \in P - \{\min P\}$, a composition $F_{P,i} = T_{i_1} \circ \cdots \circ T_{i_j}$ with $1 \leq j \leq (|P|-1)L$, such that $v_i = F_{P,i}(v_{\min P})$. We also define $F_{P,\min P}$ as the identity mapping. According to the locality of ϕ, if $\mathbf{v}' = (v_1', \ldots, v_p') \in V^p$ defines the same partition, types, and compositions, then $\mathbf{v}' \in X$. For fixed partition $\mathcal{P}$, types $(T_P)_{P\in\mathcal{P}}$, and compositions $(F_{P,i})_{i\in P\in\mathcal{P}}$, the corresponding subset X' of X is included in a (reshuffled) product Y of sets of tuples of the form $(F_{P,i}(x_{\min P}))$ for $v_{\min P} \in W_P$, and is the set of all $v \in G$ such that $B_{(|P|-1)L}(\mathbf{G}, v) = T_P$. Hence W_P is measurable and (as each $F_{P,i}$ is measurable) Y is a measurable subset of $G^{|P|}$. Of course, this product may contain tuples $\mathbf{v}$ defining another partition. A simple induction and inclusion/exclusion argument shows that X' is measurable. As X is the union of a finite number of such sets, X is measurable. □

We now relate graphings to FO-limits of bounded degree graphs. We shall make use of the following lemma which reduces a graphing to its essential support.

Lemma 3.37 (Cleaning Lemma). *Let $\mathbf{G} = (V, T_1, \ldots, T_d, \mu)$ be a graphing.*

Then there exists a subset $X \subset V$ with 0 measure such that X is globally invariant by each of the T_i and $\mathbf{G}' = (V - X, T_1, \ldots, T_d, \mu)$ is a graphing such that for every finite rooted colored graph (F, o) and integer r it holds

$$\mu(D_r(\mathbf{G}', (F, o))) = \mu(D_r(\mathbf{G}, (F, o)))$$

(which means that $\mathbf{G}'$ is equivalent to $\mathbf{G}$) and

$$D_r(\mathbf{G}', (F, o)) \neq \emptyset \quad \iff \quad \mu(D_r(\mathbf{G}', (F, o))) > 0.$$

Proof. For a fixed r, define $\mathcal{F}_r$ has the set of all (isomorphism types of) finite rooted k-edge colored graphs (F, o) with radius at most r such that $\mu(D_r(\mathbf{G}, (F, o))) = 0$. Define

$$X = \bigcup_{r\in\mathbb{N}} \bigcup_{(F,o)\in\mathcal{F}_r} D_r(\mathbf{G}, (F, o)).$$

Then $\mu(X) = 0$, as it is a countable union of 0-measure sets.

We shall now prove that X is a union of connected components of $\mathbf{G}$, and thus X is globally invariant by each of the T_i. Namely, if $x \in X$ and y is adjacent to x, then $y \in X$. Indeed: if $x \in X$ then there exists an integer r such that $\mu(D(\mathbf{G}, B_r(\mathbf{G}, x))) = 0$. But it is easily checked that

$$\mu(D(\mathbf{G}, B_{r+1}(\mathbf{G}, y))) \leq d \cdot \mu(D(\mathbf{G}, B_r(\mathbf{G}, x))).$$

Hence $y \in X$. It follows that for every $1 \leq i \leq d$ we have $T_i(X) = X$. So we can define the graphing $\mathbf{G}' = (V - X, T_1, \ldots, T_d, \mu)$.

Let (F,o) be a rooted finite colored graph. Assume there exists $x \in \mathbf{G}'$ such that $B_r(\mathbf{G}',r) \simeq (F,o)$. As X is a union of connected components, we also have $B_r(\mathbf{G},r) \simeq (F,o)$ and $x \notin X$. It follows that $\mu(D(\mathbf{G},(F,o))) > 0$ hence it holds $\mu(D_r(\mathbf{G}',(F,o))) > 0$. □

The cleaning lemma allows us a clean description of FO-limits in the bounded degree case:

THEOREM 3.38. *Let $(G_n)_{n\in\mathbb{N}}$ be a FO-convergent sequence of finite graphs with maximum degree d, with $\lim_{n\to\infty}|G_n| = \infty$. Then there exists a graphing $\mathbf{G}$, which is the disjoint union of a graphing $\mathbf{G}_0$ and a countable graph $\hat{G}$ such that*

- *The graphing $\mathbf{G}$ is a modeling* FO*-limit of the sequence $(G_n)_{n\in\mathbb{N}}$.*
- *The graphing $\mathbf{G}_0$ is a BS-limit of the sequence $(G_n)_{n\in\mathbb{N}}$ such that*

$$D_r(\mathbf{G}_0,(F,o)) \neq \emptyset \iff \mu(D_r(\mathbf{G}_0,(F,o))) > 0.$$

- *The countable graph $\hat{G}$ is an elementary limit of the sequence $(G_n)_{n\in\mathbb{N}}$.*

PROOF. Let $\mathbf{G}_0$ be a BS-limit, which has been "cleaned" using the previous lemma, and let $\hat{G}$ be an elementary limit of G. It is clear that $\mathbf{G} = \mathbf{G}_0 \cup \hat{G}$ is also a BS-limit of the sequence, so the lemma amounts in proving that $\mathbf{G}$ is elementarily equivalent to $\hat{G}$.

According to Hanf's theorem [38], it is sufficient to prove that for every integers r,t and for every rooted finite graph (F,o) (with maximum degree d) the following equality holds:

$$\min(t,|D_r(\mathbf{G},(F,o))|) = \min(t,|D_r(\hat{G},(F,o))|).$$

Assume for contradiction that this is not the case. Then $|D_r(\hat{G},(F,o))| < t$ and $D_r(\mathbf{G}_0,(F,o))$ is not empty. However, as $\mathbf{G}_0$ is clean, this implies $\mu(D_r(\mathbf{G}_0,(F,o))) = \alpha > 0$. It follows that for every sufficiently large n it holds $|D_r(G_n,(F,o))| > \alpha/2\,|G_n| > t$. Hence $|D_r(\hat{G},(F,o))| > t$, contradicting our hypothesis.

That $\mathbf{G}$ is a modeling then follows from Lemma 3.36. □

REMARK 3.39. Not every graphing with maximum degree 2 is an FO-limit modeling of a sequence of finite graphs (as it needs not be an elementary limit of finite graphs). Indeed: let $\mathbf{G}$ be a graphing that is an FO-limit modeling of the sequence of cycles. The disjoint union of $\mathbf{G}$ and a ray is a graphing $\mathbf{G}'$, which has the property that all its vertices but one have degree 2, the exceptional vertex having degree 1. As this property is not satisfied by any finite graph, $\mathbf{G}'$ is not the FO-limit of a sequence of finite graphs.

Let us finish this section by giving an interesting example, which shows that the cleaning lemma sometimes applies in a non-trivial way:

EXAMPLE 3.40. Consider the graph G_n obtained from a de Bruijn sequence of length 2^n as shown Fig 4.

It is easy to define a graphing $\mathbf{G}$, which is the limit of the sequence $(G_n)_{n\in\mathbb{N}}$: as vertex set, we consider the rectangle $[0;1) \times [0;3)$. We define a measure preserving

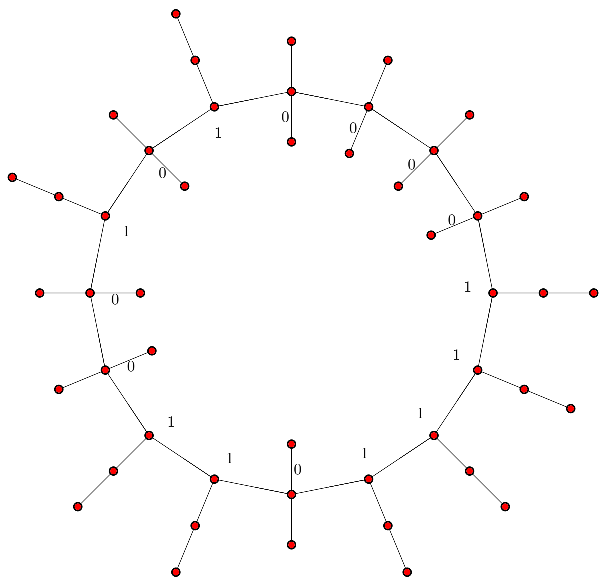

FIGURE 4. The graph G_n is constructed from a de Bruijn sequence of length 2^n.

function f and two measure preserving involutions T_1, T_2 as follows:

$$f(x,y) = \begin{cases} (2x, y/2) & \text{if } x < 1/2 \text{ and } y < 1 \\ (2x-1, (y+1)/2) & \text{if } 1/2 \leq x \text{ and } y < 1 \\ (x,y) & \text{otherwise} \end{cases}$$

$$T_1(x,y) = \begin{cases} (x, y+1) & \text{if } y < 1 \\ (x, y-1) & \text{if } 1 \leq y < 2 \\ (x,y) & \text{otherwise} \end{cases}$$

$$T_2(x,y) = \begin{cases} (x, y+1) & \text{if } x < 1/2 \text{ and } 1 \leq y < 2 \\ (x, y+2) & \text{if } 1/2 \leq x \text{ and } y < 1 \\ (x, y-1) & \text{if } x < 1/2 \text{ and } 2 \leq y \\ (x, y-2) & \text{if } 1/2 \leq x \text{ and } 2 \leq y \\ (x,y) & \text{otherwise} \end{cases}$$

Then the edges of $\mathbf{G}$ are the pairs $\{(x,y),(x',y')\}$ such that $(x,y) \neq (x',y')$ and either $(x',y') = f(x,y)$, or $(x,y) = f(x',y')$, or $(x',y') = T_1(x,y)$, or $(x',y') = T_2(x,y)$.

If one considers a random root (x,y) in $\mathbf{G}$, then the connected component of (x,y) will almost surely be a rooted line with some decoration, as expected from what is seen from a random root in a sufficiently large G_n. However, special behaviour may happen when x and y are rational. Namely, it is possible that the connected component of (x,y) becomes finite. For instance, if $x = 1/(2^n - 1)$ and $y = 2^{n-1}x$ then the orbit of (x,y) under the action of f has length n thus the connected component of (x,y) in $\mathbf{G}$ has order $3n$. Of course, such finite connected components do not appear in G_n. Hence, in order to clean $\mathbf{G}$, infinitely many components have to be removed.

Let us give a simple example exemplifying the distinction between BS and FO-convergence for graphs with bounded degree.

EXAMPLE 3.41. Let G_n denote the $n \times n$ grid. The BS-limit object is a probability distribution concentrated on the infinite grid with a specified root. A limit graphing can be described as the Lebesgue measure on $[0,1]^2$, where (x,y) is adjacent to $(x \pm \alpha \bmod 1, y \pm \alpha \bmod 1)$ for some irrational number α.

This graphing, however, is not an FO-limit of the sequence $(G_n)_{n\in\mathbb{N}}$ as every FO-limit has to contain four vertices of degree 2. An FO-limit graphing can be described as the above graphing restricted to $[0,1)^2$ (obtained by deleting all vertices with $x = 1$ or $y = 1$). One checks for instance that this graphing contains four vertices of degree 2 (the vertices (α,α), $(1-\alpha,\alpha)$, $(\alpha,1-\alpha)$, and $(1-\alpha,1-\alpha)$) and infinitely many vertices of degree 3.

We want to stress that our general and unifying approach to structural limits was not developed for its own sake and that it provided a proper setting (and, yes, encouragement) for the study of classes of sparse graphs. So far the bounded degree graphs are the only sparse class of graphs where the structural limits were constructed efficiently. (Another example of limits of sparse graphs is provided by scaling limits of transitive graphs [7] which proceeds in different direction and is not considered here.) The goal of the remaining sections of this article is to extend this to strong Borel FO-limits of rooted trees with bounded height and thus, by means of a fitting basic interpretation scheme, to graphs with bounded tree-depth (defined in [59]), or graphs with bounded SC-depth (defined in [36]).

3.3. Decomposing Sequences: the Comb Structure

The combinatorics of limits of equivalence relations (such as components) is complicated. We start this analysis by considering the combinatorics of "large" equivalence classes. This leads to the notion of spectrum, which will be analyzed in this section.

3.3.1. Spectrum of a First-order Equivalence Relation.

DEFINITION 3.42 (ϖ-spectrum). Let $\mathbf{A}$ be a λ-modeling (with measure $\nu_{\mathbf{A}}$), and let $\varpi \in \mathrm{FO}_2(\lambda)$ be a formula expressing a component relation on $\mathbf{A}$ (see Definition 2.22). Let $\{C_i : i \in \Gamma\}$ be the set of all the ϖ-equivalence classes of A, and let Γ_+ be the (at most countable) subset of Γ of the indexes i such that $\nu_{\mathbf{A}}(C_i) > 0$.

The *ϖ-spectrum* $\mathrm{Sp}_\varpi(\mathbf{A})$ of $\mathbf{A}$ is the (at most countable) sequence of the values $\nu_{\mathbf{A}}^p(C_i)$ (for $i \in \Gamma_+$) ordered in non-increasing order.

LEMMA 3.43. *For $k \in \mathbb{N}$, let $\varpi^{(k)}$ be the formula $\bigwedge_{i=1}^k \varpi(x_i, x_{i+1})$. Then it holds*

$$\sum_{i\in\Gamma_+} \nu_{\mathbf{A}}(C_i)^{k+1} = \langle \varpi^{(k)}, \mathbf{A}\rangle.$$

PROOF. Let $k \in \mathbb{N}$. Define

$$D_{k+1} = \{(x_1, \dots, x_{k+1}) \in A^{k+1} :\ A \models \varpi^k(x_1, \dots, x_{k+1})\}.$$

According to Lemma 3.3, each C_i is measurable, thus $\bigcup_{i\in\Gamma_+} C_i$ is measurable and so is $R = A \setminus \bigcup_{i\in\Gamma_+} C_i$.

Considering the indicator function $\mathbf{1}_{D_{k+1}\cap R^{k+1}}$ of $D_{k+1} \cap R^{k+1}$ and applying Fubini's theorem, we get

$$\int_{A^{k+1}} \mathbf{1}_{D_{k+1}\cap R^{k+1}}\, \mathrm{d}\nu_{\mathbf{A}}^{k+1} = \int \cdots \int \mathbf{1}_R(x_1, \dots, x_{k+1})\ \mathrm{d}\nu_{\mathbf{A}}(x_1, \dots, \mathrm{d}\nu_{\mathbf{A}}(x_{k+1}) = 0.$$

as for every fixed $a_1, \dots, a_k$ (with $a_1 \in C_\alpha$, for some $\alpha \in \Gamma \setminus \Gamma_+$) we have

$$0 \leq \int \mathbf{1}_R(a_1, \dots, a_k, x_{k+1})\, \mathrm{d}\nu_{\mathbf{A}}(x_{k+1}) \leq \nu_{\mathbf{A}}(C_\alpha) = 0.$$

It follows (by countable additivity) that

$$\langle \varpi^{(k)}, \mathbf{A}\rangle = \nu_{\mathbf{A}}^{k+1}(D_{k+1}) = \nu_{\mathbf{A}}^{k+1}(\bigcup_{i\in\Gamma_+} C_i^{k+1}) = \sum_{i\in\Gamma_+} \nu_{\mathbf{A}}(C_i)^{k+1}.$$

$\square$

It follows from Lemma 3.43 that the spectrum $\mathrm{Sp}_\varpi(\mathbf{A})$ is computable from the sequence of (non-increasing) values $(\langle \varpi^{(k)}, \mathbf{A}\rangle)_{k\in\mathbb{N}}$.

We assume that every finite sequence $\mathbf{x} = (x_1, \dots, x_n)$ of positive reals is implicitly embedded in an infinite sequence by defining $x_i = 0$ for $i > n$. Recall the usual ℓ_k norms:

$$\|\mathbf{x}\|_k = \Big(\sum_i |x_i|^k\Big)^{1/k}.$$

Hence above equations rewrite as

$$\|\mathrm{Sp}_\varpi(\mathbf{A})\|_{k+1} = \langle \varpi^{(k)}, \mathbf{A}\rangle^{1/(k+1)} \tag{3.1}$$

We shall prove that the spectrum is, in a certain sense, defined by a continuous function. We need the following technical lemma.

LEMMA 3.44. *For each $n \in \mathbb{N}$, let $\mathbf{a}_n = (a_{n,i})_{i\in\mathbb{N}}$ be a non-increasing sequence of positive real numbers with bounded sum (i.e $\|\mathbf{a}_n\|_1 < \infty$ for every $n \in \mathbb{N}$).*

Assume that for every integer $k \geq 1$ the limit $s_k = \lim_{n\to\infty} \|\mathbf{a}_n\|_k$ exists.

Then $(\mathbf{a}_n)_{n\in\mathbb{N}}$ converges in the space c_0 of all sequences converging to zero (with norm $\|\cdot\|_\infty$).

PROOF. We first prove that the sequences converge pointwise, that is that there exists a sequence $\mathbf{x} = (x_i)_{i\in\mathbb{N}}$ such that for every $i \in \mathbb{N}$ it holds

$$x_i = \lim_{n\to\infty} a_{n,i}.$$

For every $\epsilon > 0$, if $s_k < \epsilon$ then $a_{n,1} < 2\epsilon$ for all sufficiently large values of n. Thus if $s_k = 0$ for some k, the limit $\lim_{n\to} a_{n,i}$ exists for every i and is null. Thus, we can assume that s_k is strictly positive for every $k \in \mathbb{N}$.

Fix $k \in \mathbb{N}$. There exists $N \in \mathbb{N}$ such that for every $n \geq N$ it holds $|s_k^k - \|\mathbf{a}_n\|_k^k| < s_k^k/k$. As $(a_{n,i})_{i\in\mathbb{N}}$ is a non-increasing sequence of positive real numbers, for every $n \neq N$ it holds

$$a_{n,1}^k \leq \|\mathbf{a}_n\|_k^k < s_k^k(1 + 1/k)$$

and

$$a_{n,1}^{k-1} \geq \|\mathbf{a}_n\|_k^k > s_k^k(1-1/k)$$

Hence

$$\log s_k + \frac{\log(1+1/k)}{k} \geq \log a_{n,1} \geq \big(1+\frac{1}{k-1}\big)\big(\log s_k + \frac{\log(1-1/k)}{k}\big)$$

Thus $x_1 = \lim_{n\to\infty} a_{n,1}$ exists and $x_1 = \lim_{k\to\infty} s_k$. Inductively, we get that for each $i \in \mathbb{N}$, the limit $x_i = \lim_{n\to\infty} a_{n,i}$ exists and that

$$x_i = \lim_{k\to\infty} (s_k^k - \sum_{j<i} x_j^k)^{1/k}.$$

We now prove that the converge is uniform, that is that for every $\epsilon > 0$ there exists N such that for every $n \geq N$ it holds

$$\|\mathbf{x} - \mathbf{a}_n\|_\infty < \epsilon.$$

As $\mathbf{a}_n \in \ell_1$ and $\|\mathbf{a}_n\|_1$ converges there exists M such that $\|\mathbf{a}_n\|_1 \leq M$ for every $n \in \mathbb{N}$. Let $\epsilon > 0$. Let $A = \min\{i : x_i \leq \epsilon/3\}$. (Note that $A \leq 3M/\epsilon$.) There exists N such that for every $n \geq N$ it holds $\sup_{i\leq A} |x_i - a_{n,i}| < \epsilon/3$. Moreover, for every $i > A$ it holds

$$0 \leq a_{n,i} \leq a_{n,A} < x_A + \epsilon/3 < 2\epsilon/3.$$

As $0 \leq x_i \leq \epsilon/3$ for every $i > A$ it holds

$$|x_i - a_{n,i}| < \epsilon$$

for every $i > A$ (hence for every i). Thus $(\mathbf{a}_n)_{n\in\mathbb{N}}$ converges in ℓ_∞. As obviously each $\mathbf{a}_n$ has 0 limit, $(\mathbf{a}_n)_{n\in\mathbb{N}}$ converges in c_0. □

LEMMA 3.45. *Let λ be a signature. The mapping $\mathbf{A} \mapsto \mathrm{Sp}_\varpi(\mathbf{A})$ is a continuous mapping from the space of λ-modelings with a component relation ϖ (with the topology of $\mathrm{FO}^{\mathrm{local}}(\lambda)$-convergence) to the space c_0 of all sequences converging to zero (with $\|\cdot\|_\infty$ norm).*

PROOF. Assume $\mathbf{A}_n$ is an $\mathrm{FO}^{\mathrm{local}}(\lambda)$-convergent sequence of λ-modelings.

Let $(\mathsf{sp}_{n,1}, \dots, \mathsf{sp}_{n,i}, \dots)$ be the ϖ-spectrum of $\mathbf{A}_n$ (extended by zero values if finite), and let $\mathbf{a}_n = (a_{n,i})_{i\in\mathbb{N}}$ be the sequence defined by $a_{n,i} = \mathsf{sp}_{n,i}^2$. Then for every integer $k \geq 1$ it holds

$$\|a_n\|_k = \|\mathrm{Sp}_\varpi(\mathbf{A}_n)\|_{2k}^2 = \langle \varpi^{(2k-1)}, \mathbf{A}_n\rangle^{1/k}.$$

Hence $s_k = \lim_{n\to\infty} \|a_n\|_k$ exists. According to Lemma 3.44, $(\mathbf{a}_n)_{n\in\mathbb{N}}$ converges in c_0, thus so does $(\mathrm{Sp}_\varpi(\mathbf{A}_n))_{n\in\mathbb{N}}$. □

DEFINITION 3.46. Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be a sequence of finite λ-structures. Let ϖ be a component relation, and for simplicity assume $\varpi \in \lambda$. In the following, we assume that ϖ-spectra are extended to infinite sequences by adding zeros if necessary.

- The sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is *ϖ-nice* if $\mathrm{Sp}_\varpi(\mathbf{A}_n)$ converges pointwise;
- The *limit ϖ-spectrum* of a ϖ-nice sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is the pointwise limit of $\mathrm{Sp}_\varpi(\mathbf{A}_n)$;
- the *ϖ-support* is the set I of the indexes i for which the limit ϖ-spectrum is non-zero;

- the sequence has *full* ϖ-spectrum if, for every index i not in the ϖ-support, there is some N such that the ith value of $\mathrm{Sp}_\varpi(\mathbf{A}_n)$ is zero for every $n > N$.

As proved in Lemma 3.45, every $\mathrm{FO}^{\mathrm{local}}$-convergent sequence is ϖ-nice.

LEMMA 3.47. *Let $(\mathbf{A}_n)$ be a ϖ-nice sequence of λ-structures with empty ϖ-support.*

Then the following conditions are equivalent:

(1) the sequence $(\mathbf{A}_n)$ is $\mathrm{FO}^{\mathrm{local}}$-convergent;
(2) the sequence $(\mathbf{A}_n)$ is $\mathrm{FO}_1^{\mathrm{local}}$-convergent.

Moreover, for every ϖ-local formula ϕ with $p > 1$ free variables it holds

$$\lim_{n\to\infty} \langle \phi, \mathbf{A}_n \rangle = 0.$$

PROOF. $\mathrm{FO}^{\mathrm{local}}$-convergence obviously implies $\mathrm{FO}_1^{\mathrm{local}}$-convergence. So, assume that $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}_1^{\mathrm{local}}$-convergent, and let ϕ be a ϖ-local first-order formula with $p > 1$ free variables. For $n \in \mathbb{N}$, let $\mathbf{B}_{n,i}$ ($i \in \Gamma_n$) denote the connected components of $\mathbf{A}_n$. As $(\mathbf{A}_n)$ is ϖ-nice and has empty ϖ-support, there exists for every $\epsilon > 0$ an integer N such that for $n > N$ and every $i \in \Gamma_n$ it holds $|B_{n,i}| < \epsilon |A_n|$. Then, according to Corollary 3.3, for $n > N$

$$\begin{aligned}
\langle \phi, \mathbf{A}_n \rangle &= \sum_{i\in\Gamma_n} \left(\frac{|B_{n,i}|}{|A_n|}\right)^p \langle \phi, \mathbf{B}_{n,i} \rangle \\
&\le \sum_{i\in\Gamma_n} \left(\frac{|B_{n,i}|}{|A_n|}\right)^p \\
&< \sum_{i\in\Gamma_n} \frac{|B_{n,i}|}{|A_n|} \epsilon^{p-1} = \epsilon^{p-1}
\end{aligned}$$

Hence $\langle \phi, \mathbf{A}_n \rangle$ converges (to 0) as n grows to infinity. It follows that $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent, according to Theorem 3.17. □

LEMMA 3.48. *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be an $\mathrm{FO}^{\mathrm{local}}$-convergent sequence of finite λ-structures, with component relation $\varpi \in \lambda$ and limit ϖ-spectrum $(\mathsf{sp}_i)_{i\in I}$. For $n \in \mathbb{N}$, let $\mathbf{B}_{n,i}$ be the connected components of $\mathbf{A}_n$ order in non-decreasing order (with $\mathbf{B}_{n,i}$ empty if i is greater than the number of connected components of $\mathbf{A}_n$). Let $a \le b$ be the first and last occurrence of $\mathsf{sp}_a = \mathsf{sp}_b$ in the ϖ-spectrum and let $\mathbf{A}'_n$ be the union of all the $\mathbf{B}_{n,i}$ for $a \le i \le b$.*

Then $(\mathbf{A}'_n)_{n\in\mathbb{N}}$ is FO-convergent if $\mathsf{sp}_a > 0$ and $\mathrm{FO}^{\mathrm{local}}$-convergent if $\mathsf{sp}_a = 0$.

Assume moreover that $(\mathbf{A}_n)_{n\in\mathbb{N}}$ has a modeling $\mathrm{FO}^{\mathrm{local}}$-limit $\mathbf{L}$. Let $\mathbf{L}'$ be the union of the connected components $\mathbf{L}_i$ of $\mathbf{L}$ with $\nu_{\mathbf{L}}(L_i) = \mathsf{sp}_a$. Equip $\mathbf{L}'$ with the σ-algebra $\Sigma_{\mathbf{L}'}$ which is the restriction of $\Sigma_{\mathbf{L}}$ to L' and the probability measure $\nu_{\mathbf{L}'}$ defined by $\nu_{\mathbf{L}'}(X) = \nu_{\mathbf{L}}(X)/\nu_{\mathbf{L}}(L')$ (for $X \in \Sigma_{\mathbf{L}'}$).

Then $\mathbf{L}'$ is a modeling FO-limit of $(\mathbf{A}'_n)_{n\in\mathbb{N}}$ if $\mathsf{sp}_a > 0$ and a modeling $\mathrm{FO}^{\mathrm{local}}$-limit of $(\mathbf{A}'_n)_{n\in\mathbb{N}}$ if $\mathsf{sp}_a = 0$.

PROOF. Extend the sequence sp to the null index by defining $\mathsf{sp}_0 = 2$. Let $r = \min(\mathsf{sp}_{a-1}/\mathsf{sp}_a, \mathsf{sp}_b/\mathsf{sp}_{b+1})$ (if $\mathsf{sp}_{b+1} = 0$ simply define $r = \mathsf{sp}_{a-1}/\mathsf{sp}_a$). Notice that

$r > 1$. Let ϕ be a ϖ-local formula with p free variables. According to Corollary 3.3 it holds

$$\langle \phi, \mathbf{A}_n \rangle = \sum_i \left(\frac{|B_{n,i}|}{|A_n|} \right)^p \langle \phi, \mathbf{B}_{n,i} \rangle.$$

In particular, it holds

$$\langle \varpi^{(p)}, \mathbf{A}_n \rangle = \sum_i \left(\frac{|B_{n,i}|}{|A_n|} \right)^p.$$

Let $\alpha > 1/(1 - r^p)$. Define

$$w_{n,i} = \left(\frac{|B_{n,i}|}{|A_n|} \right)^p (\alpha + \langle \phi, \mathbf{B}_{n,i} \rangle).$$

¿From the definition of r it follows that for each $n \in \mathbb{N}$, $w_{n,i} > w_{n,j}$ if $i < a$ and $j \geq a$ or $i \leq b$ and $j > b$. Let $\sigma \in S_\omega$ be a permutation of $\mathbb{N}$ with finite support, such that $a_{n,i} = w_{n,\sigma(i)}$ is non-increasing. It holds

$$\sum_i a_{n,i} = \sum_i w_{n,i} = \alpha \langle \varpi^{(p)}, \mathbf{A}_n \rangle + \langle \phi, \mathbf{A}_n \rangle.$$

Hence

$$\lim_{n \to \infty} \sum_i a_{n,i}^p$$

exists. According to Lemma 3.44 it follows that for every $i \in \mathbb{N}$ the limit $\lim_{n \to \infty} a_{n,i}$ exists. Moreover, as σ globally preserves the set $\{a, \ldots, b\}$ it follows that the limit

$$d = \lim_{n \to \infty} \sum_{i=a}^{b} \left(\frac{|B_{n,i}|}{|A_n|} \right)^p (\alpha + \langle \phi, \mathbf{B}_{n,i} \rangle)$$

exists. As for every $i \in \{a, \ldots, b\}$ it holds $\lim_{n \to \infty} |B_{n,i}|/|A_n| = \mathsf{sp}_a$ and as $\langle \phi, \mathbf{A}'_n \rangle = \sum_{i=a}^{b} (|B_{n,i}|/|A_n|)^p \langle \phi, \mathbf{B}_{n,i} \rangle$ we deduce

$$\lim_{n \to \infty} \langle \phi, \mathbf{A}'_n \rangle = d - (b - a + 1)\alpha.$$

Hence $\lim_{n \to \infty} \langle \phi, \mathbf{A}'_n \rangle$ exists for every ϖ-local formula and, according to Theorem 3.17, the sequence $(\mathbf{A}'_n)_{n \in \mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent.

Assume $\mathsf{sp}_a > 0$. Let $N = b - a + 1$. To each sentence θ we associate the formula $\widetilde{\theta} \in \mathrm{FO}_N^{\mathrm{local}}$ that asserts that the substructure induced by the closed neighborhood of $x_1, \ldots, x_N$ satisfies θ and that $x_1, \ldots, x_N$ are pairwise distinct and non-adjacent. For sufficiently large n, the structure $\mathbf{A}'_n$ has exactly N connected components. It is easily checked that if $\mathbf{A}'_n$ does not satisfy θ then $\langle \widetilde{\theta}, \mathbf{A}'_n \rangle = 0$, although if $\mathbf{A}'_n$ does satisfy θ then

$$\langle \widetilde{\theta}, \mathbf{A}'_n \rangle \geq \left(\frac{\min_{a \leq i \leq b} |B_{n,i}|}{\sum_{i=a}^{b} |B_{n,i}|} \right)^N,$$

hence $\langle \widetilde{\theta}, \mathbf{A}'_n \rangle > (2N)^{-N}$ for all sufficiently large n. As $\langle \widetilde{\theta}, \mathbf{A}'_n \rangle$ converges for every sentence θ, we deduce that the sequence $(\mathbf{A}'_n)_{n \in \mathbb{N}}$ is elementarily convergent. According to Theorem 2.21, the sequence $(\mathbf{A}'_n)_{n \in \mathbb{N}}$ is thus FO-convergent.

Now assume that $(\mathbf{A}_n)_{n\in\mathbb{N}}$ has a modeling $\mathrm{FO}^{\mathrm{local}}$-limit $\mathbf{L}$. First note that L_i being an equivalence class of ϖ it holds $L_i \in \Sigma_{\mathbf{L}}$, hence $L' \in \Sigma_{\mathbf{L}}$ and $\nu_{\mathbf{L}}(L')$ is well defined. For every ϖ-local formula $\phi \in \mathrm{FO}_p(\lambda)$ it holds, according to Lemma 3.16:

$$\begin{aligned}\langle \phi, \mathbf{L}'\rangle &= \sum_{i=a}^{b} \nu_{\mathbf{L}'}(L_i)^p \langle \phi, \mathbf{L}_i\rangle \\ &= \frac{1}{\nu_{\mathbf{L}}(L')^p} \sum_{i=a}^{b} \nu_{\mathbf{L}}(L_i)^p \langle \phi, \mathbf{L}_i\rangle\end{aligned}$$

We deduce that

$$\langle \phi, \mathbf{L}'\rangle = \lim_{n\to\infty} \langle \phi, \mathbf{A}'_n\rangle.$$

According to Theorem 3.17, it follows that the same equality holds for every $\phi \in \mathrm{FO}^{\mathrm{local}}(\lambda)$ hence $\mathbf{L}'$ is a modeling $\mathrm{FO}^{\mathrm{local}}$-limit of the sequence $(\mathbf{A}'_n)_{n\in\mathbb{N}}$.

As above, for $\mathsf{sp}_a > 0$, if $\mathbf{L}'$ is a modeling $\mathrm{FO}^{\mathrm{local}}$-limit of $(\mathbf{A}'_n)_{n\in\mathbb{N}}$ then it is a modeling FO-limit.

□

LEMMA 3.49. *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be an* FO*-convergent sequence of finite λ-structures, with component relation ϖ (expressing usual notion of connected components). Assume all the $\mathbf{A}_n$ have at most k connected components. Denote by $\mathbf{B}_{n,1}, \dots, \mathbf{B}_{n,k}$ these components (adding empty λ-structures if necessary).*

Assume that for each $1 \le i \le k$ it holds $\lim_{n\to\infty} |B_{n,i}|/|A_n| = 1/k$.

Then there exists a sequence $(\sigma_n)_{n\in\mathbb{N}}$ of permutations of $[k]$ such that for each $1 \le i \le k$ the sequence $(\mathbf{B}_{n,\sigma_n(i)})_{n\in\mathbb{N}}$ is FO*-convergent.*

PROOF. To a formula $\phi \in \mathrm{FO}_p(\lambda)$ we associate the ϖ-local formula $\widetilde{\phi} \in \mathrm{FO}_p^{\mathrm{local}}(\lambda)$ asserting that all the free variables are ϖ-adjacent and that their closed neighborhood (that is their connected component) satisfies ϕ. Then essentially the same proof as above allows to refine $\mathbf{A}_n$ into sequences such that $\langle \phi, \mathbf{A}'_{n,i}\rangle$ is constant on the connected components of each of the $\mathbf{A}'_n$. Considering formulas allowing to split at least one of the sequences, we repeat this process (at most $k-1$ times) until each $\mathbf{A}'_{n,i}$ contains equivalent connected components. Then, $\mathbf{A}'_{n,i}$ can be split into connected components in an arbitrary order, thus obtaining the sequences $\mathbf{B}_{n,i}$. □

So we have proved that a FO-convergent can be decomposed by isolines of the ϖ-spectrum. In the next sections, we shall investigate how to refine further.

3.3.2. Sequences with Finite Spectrum. For every ϖ-nice sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ with finite support I, we define the *residue* $\mathbf{R}_n$ of $\mathbf{A}_n$ as the union of the connected components $\mathbf{B}_{n,i}$ of $\mathbf{A}_n$ such that $i \notin I$.

When one considers an $\mathrm{FO}^{\mathrm{local}}$-convergent sequence $(\mathbf{A}_n)$ with a finite support then the sequence of the residues forms a sequence which is either $\mathrm{FO}^{\mathrm{local}}$-convergent or "negligible" in the sense that $\lim_{n\to\infty} |R_n|/|A_n| = 0$. This is formulated as follows:

LEMMA 3.50. *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be a sequence of λ-structures with component relation ϖ. For each $n \in \mathbb{N}$ and $i \in \mathbb{N}$, let $\mathbf{B}_{n,i}$ be the i-th largest connected component of $\mathbf{A}_n$.*

Assume that $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *is* $\mathrm{FO}^{\mathrm{local}}$*-convergent and has finite spectrum* $(\mathsf{sp}_i)_{i\in I}$*. Let* $\mathbf{R}_n$ *be the residue of* $\mathbf{A}_n$*.*

Then $\mathsf{sp}' = \lim_{n\to\infty}|R_n|/|A_n|$ *exists and either* $\mathsf{sp}'=0$ *or* $(\mathbf{R}_n)_{n\in\mathbb{N}}$ *is* $\mathrm{FO}^{\mathrm{local}}$*-convergent.*

Proof. Clearly, $\mathsf{sp}' = 1-\sum_i \mathsf{sp}_i$. Assume $\mathsf{sp}'>0$. First notice that for every $\epsilon>0$ there exists N such that for every $i>N$, the λ-structure $\mathbf{R}_n$ has no connected component of size at least $\epsilon/2\mathsf{sp}'|A_n|$ and $\mathbf{R}_n$ has order at least $\mathsf{sp}'/2|A_n|$. Hence, for every $i>N$, the λ-structure $\mathbf{R}_n$ has no connected component of size at least $\epsilon|R_n|$. According to Lemma 3.47, proving that $(\mathbf{R}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent reduces to proving that $(\mathbf{R}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}_1^{\mathrm{local}}$-convergent.

Let $\phi\in\mathrm{FO}_1^{\mathrm{local}}$. We group the λ-structures $\mathbf{B}_{n,i}$ (for $i\in I$) by values of sp_i as $\mathbf{A}'_{n,1},\dots,\mathbf{A}'_{n,q}$. Denote by c_j the common value of sp_i for the connected components $\mathbf{B}_{n,i}$ in $\mathbf{A}'_{n,j}$. According to Corollary 3.3 it holds (as ϕ is clearly ϖ-local):

$$\begin{aligned}\langle\phi,\mathbf{A}_n\rangle &= \sum_i \frac{|B_{n,i}|}{|A_n|}\langle\phi,\mathbf{B}_{n,i}\rangle\\ &= \sum_{i\in I}\frac{|B_{n,i}|}{|A_n|}\langle\phi,\mathbf{B}_{n,i}\rangle + \sum_{i\notin I}\frac{|B_{n,i}|}{|A_n|}\langle\phi,\mathbf{B}_{n,i}\rangle\\ &= \sum_{j=1}^{q}\frac{|A'_{n,j}|}{|A_n|}\langle\phi,\mathbf{A}'_{n,j}\rangle + \frac{|R_n|}{|A_n|}\langle\phi,\mathbf{R}_n\rangle\end{aligned}$$

According to Lemma 3.48, each sequence $(\mathbf{A}'_{n,j})_{n\in\mathbb{N}}$ is FO-convergent. Hence the limit $\lim_{n\to\infty}\langle\phi,\mathbf{R}_n\rangle$ exists and we have

$$\lim_{n\to\infty}\langle\phi,\mathbf{R}_n\rangle = \frac{1}{\mathsf{sp}'}\left(\lim_{n\to\infty}\langle\phi,\mathbf{A}_n\rangle - \sum_{j=1}^{q} c_j \lim_{n\to\infty}\langle\phi,\mathbf{A}'_{n,j}\rangle\right).$$

It follows that the sequence $(\mathbf{R}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent. $\square$

The following result finally determines the structure of converging sequences of (disconnected) λ-structures with finite support. This structure is called *comb structure*, see Fig 5.

Theorem 3.51 (Comb structure for λ-structure sequences with finite spectrum). *Let* $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *be an* $\mathrm{FO}^{\mathrm{local}}$*-convergent sequence of finite* λ*-structures with component relation* ϖ *and finite spectrum* $(\mathsf{sp}_i)_{i\in I}$*. Let* $\mathbf{R}_n$ *be the residue of* $\mathbf{A}_n$*.*

Then there exists, for each $n\in\mathbb{N}$*, a permutation* $f_n: I\to I$ *such that it holds*

- $\lim_{n\to\infty}\max_{i\notin I}|B_{n,i}|/|A_n| = 0$*;*
- $\lim_{n\to\infty}|R_n|/|A_n|$ *exists;*
- *for every* $i\in I$*, the sequence* $(B_{n,f_n(i)})_{n\in\mathbb{N}}$ *is* FO*-convergent and* $\lim_{n\to\infty}|B_{n,f_n(i)}|/|A_n| = \mathsf{sp}_i$*;*
- *either* $\lim_{n\to\infty}|R_n|/|A_n| = 0$*, or the sequence* $(\mathbf{R}_n)_{n\in\mathbb{N}}$ *is* $\mathrm{FO}^{\mathrm{local}}$*-convergent.*

Moreover, if $(\mathbf{A}_n)_{n\in\mathbb{N}}$ *is* FO*-convergent then* $(\mathbf{R}_n)_{n\in\mathbb{N}}$ *is elementary-convergent.*

Proof. This lemma is a direct consequence of Lemmas 3.48, 3.49 and 3.50, except that we still have to prove FO-convergence of $(\mathbf{R}_n)_{n\in\mathbb{N}}$ in the case where

$(\mathbf{A}_n)_{n\in\mathbb{N}}$ is FO-convergent. As I is finite, the elementary convergence of $(\mathbf{R}_n)_{n\in\mathbb{N}}$ easily follows from the one of $(\mathbf{A}_n)$ and the one of the $(\mathbf{B}_{n,f_n(i)})$ for $i \in I$. $\square$

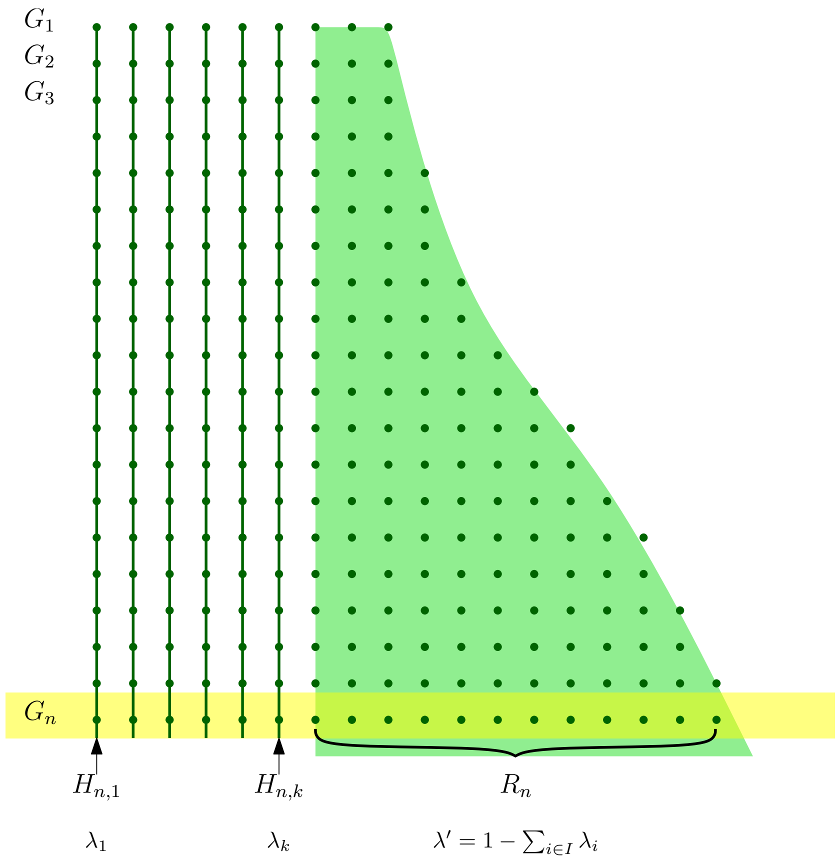

FIGURE 5. Illustration of the Comb structure for sequences with finite support

3.3.3. Sequences with Infinite Spectrum. Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be a ϖ-nice sequence with infinite spectrum (and support $I = \mathbb{N}$). In such a case, the notion of a residue becomes more tricky and will need some technical definitions. Before this, let us take the time to give an example illustrating the difficulty of the determination of the residue $\mathbf{R}_n$ in the comb structure of sequences with infinite spectrum.

EXAMPLE 3.52. Consider the sequence $(G_n)_{n\in\mathbb{N}}$ where G_n is the union of 2^n stars $H_{n,1}, \dots, H_{n,2^n}$, where the i-th star $H_{n,i}$ has order $2^{2^n}(2^{-i} + 2^{-n})/2$. Then it holds

$$\mathsf{sp}_i = \lim_{n\to\infty} |H_{n,i}|/|G_n| = 2^{-(i+1)}$$

hence $\sum_i \mathsf{sp}_i = 1/2$ thus the residue asymptotically should contain half of the vertices of G_n! An FO-limit of this sequence is shown Fig. 6.

This example is not isolated. In fact it is quite frequent in many of its variants. To decompose such examples we need a convenient separation. This is provided by the notion of clip.

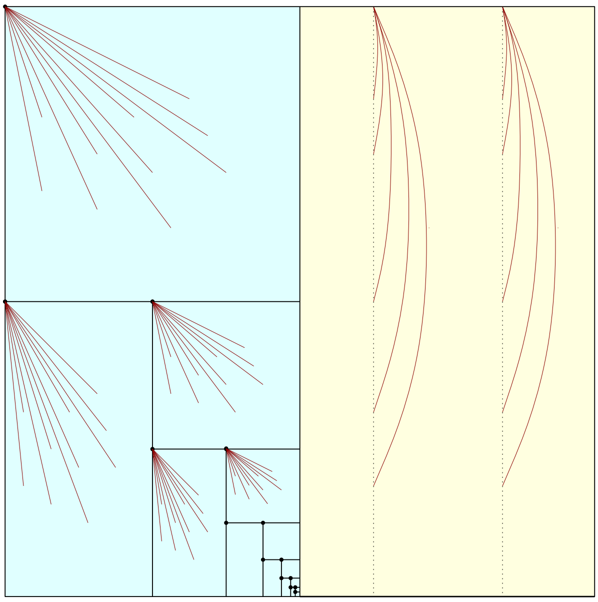

FIGURE 6. An FO-local limit. On the left side, each rectangle correspond to a star with the upper left point as its center; on the right side, each vertical line is a star with the upper point as its center.

DEFINITION 3.53.
- A *clip* of a ϖ-nice sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ with support $\mathbb{N}$ is a non-decreasing function $C:\mathbb{N}\to\mathbb{N}$ such that $\lim_{n\to\infty} C(n)=\infty$ and
$$\forall n'\geq n\quad \sum_{i=1}^{C(n)}\left|\frac{|B_{n',i}|}{|A_{n'}|}-\mathsf{sp}_i\right|\leq\sum_{i>C(n)}\mathsf{sp}_i$$
- The *residue* $\mathbf{R}_n$ of $\mathbf{A}_n$ with respect to a clip $C(n)$ is the disjoint union of the $\mathbf{B}_{n,i}$ for $i>C(n)$.

PROPOSITION 3.2. *Every ϖ-nice sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$ with infinite support has a clip C_0, which is defined by*
$$C_0(n)=\sup\left\{M,\quad(\forall n'\geq n)\ \sum_{i=1}^{M}\left|\frac{|B_{n',i}|}{|A_{n'}|}-\mathsf{sp}_i\right|\leq\sum_{i>M}\mathsf{sp}_i\right\}.$$
Moreover, $\lim_{n\to} C_0(n)=\infty$ and a non decreasing function C is a clip of $(\mathbf{A}_n)_{n\in\mathbb{N}}$ if and only if $C\leq C_0$ and $\lim_{n\to} C(n)=\infty$.

PROOF. Indeed, for each $n\in\mathbb{N}$, the value $z_l(M)=\sup_{n'\geq n}\sum_{i=1}^{M}\left|\frac{|B_{n',i}|}{|A_{n'}|}-\mathsf{sp}_i\right|$ is non-decreasing function of C with $z_l(0)=0$, and $z_r(M)=\sum_{i>M}\mathsf{sp}_i$ is a decreasing function of C with $z_r(0)=\sum_i\mathsf{sp}_i>0$ hence C_0 is well defined. Moreover, for every integer M, let $\alpha=\sum_{i>M}\mathsf{sp}_i>0$. Then, as $\lim_{n\to\infty}|B_{n',i}|/|A_{n'}|=\mathsf{sp}_i$ there exists N such that for every $n'\geq N$ and every $1\leq i\leq M$ it holds $||B_{n',i}|/|A_{n'}|-\mathsf{sp}_i|\leq\alpha/M$ thus for every $n'\geq N$ it holds
$$\sum_{i=1}^{M}\left|\frac{|B_{n',i}|}{|A_{n'}|}-\mathsf{sp}_i\right|\leq\alpha=\sum_{i>M}\mathsf{sp}_i.$$
It follows that $C_0(N)\geq M$. Hence $\lim_{n\to\infty}C_0(n)=\infty$.

That a non decreasing function C is a clip of $(\mathbf{A}_n)_{n\in\mathbb{N}}$ if and only if $C \leq C_0$ and $\lim_{n\to} C(n) = \infty$ follows directly from the definition. □

LEMMA 3.54. *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be a ϖ-nice sequence with support $\mathbb{N}$, and let C be a clip of $(\mathbf{A}_n)_{n\in\mathbb{N}}$.*

Then the limit $\mathsf{sp}' = \lim_{n\to\infty}\frac{|R_n|}{|A_n|}$ exists and $\mathsf{sp}' = 1 - \sum_i \mathsf{sp}_i$.

PROOF. As C is a clip, it holds for every $n \in \mathbb{N}$

$$\sum_i \mathsf{sp}_i - 2\sum_{i>C(n)} \mathsf{sp}_i \leq \sum_{i=1}^{C(n)} \frac{|B_{n,i}|}{|A_n|} \leq \sum_i \mathsf{sp}_i.$$

Also, for every $\epsilon > 0$ there exists n such that $|\sum_{i=1}^{C(n)} \mathsf{sp}_i - \sum_i \mathsf{sp}_i| < \epsilon$, that is: $\sum_{i>C(n)} \mathsf{sp}_i < \epsilon$. It follows that

$$\lim_{n\to\infty} \sum_{i=1}^{C(n)} \frac{|B_{n,i}|}{|A_n|} = \sum_i \mathsf{sp}_i.$$

Hence the limit $\mathsf{sp}' = \lim_{n\to\infty}\frac{|R_n|}{|A_n|}$ exists and $\mathsf{sp}' = 1 - \sum_i \mathsf{sp}_i$. □

LEMMA 3.55. *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be a sequence of λ-structures with component relation ϖ. For each $n \in \mathbb{N}$ and $i \in \mathbb{N}$, let $\mathbf{B}_{n,i}$ be the i-th largest connected component of $\mathbf{A}_n$ (if i is at most equal to the number of connected components of $\mathbf{A}_n$, the empty λ-structure otherwise).*

Assume that $(\mathbf{A}_n)_{n\in\mathbb{N}}$ is FO*-convergent.*

Let $C : \mathbb{N} \to \mathbb{N}$ be a clip of $(\mathbf{A}_n)_{n\in\mathbb{N}}$, and let $\mathbf{R}_n$ be the residue of $\mathbf{A}_n$ with respect to C.

Let $\mathsf{sp}' = \lim_{n\to\infty} |R_n|/|A_n|$. Then either $\mathsf{sp}' = 0$ or $(\mathbf{R}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent.

PROOF. According to Lemma 3.54, $\lim_{n\to\infty} |R_n|/|A_n|$ exists and $\mathsf{sp}' = 1 - \sum_i \mathsf{sp}_i$. Assume $\mathsf{sp}' > 0$. First notice that for every $\epsilon > 0$ there exists N such that for every $i > N$, the λ-structure $\mathbf{R}_n$ has no connected component of size at least $\epsilon/2\mathsf{sp}'|\mathbf{A}_n|$ and $\mathbf{R}_n$ has order at least $\mathsf{sp}'/2|A_n|$. Hence, for every $i > N$, the λ-structure $\mathbf{R}_n$ has no connected component of size at least $\epsilon|R_n|$. According to Lemma 3.47, proving that $(\mathbf{R}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent reduces to proving that $(\mathbf{R}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}_1^{\mathrm{local}}$-convergent.

Let $\phi \in \mathrm{FO}_1^{\mathrm{local}}$ (thus ϕ is ϖ-local). Let $\epsilon > 0$. There exists $k \in \mathbb{N}$ such that $\sum_{i\leq k} \mathsf{sp}_i > 1 - \mathsf{sp}' - \epsilon/3$ and such that $\mathsf{sp}_{k+1} < \mathsf{sp}_k$. We group the λ-structures $\mathbf{B}_{n,i}$ (for $1 \leq i \leq k$) by values of sp_i as $\mathbf{A}'_{n,1}, \dots, \mathbf{A}'_{n,q}$. Denote by c_j the common value of sp_i for the connected components $\mathbf{B}_{n,i}$ in $\mathbf{A}'_{n,j}$. According to Lemma 3.48, each sequence $(\mathbf{A}'_{n,i})_{n\in\mathbb{N}}$ is FO-convergent. Define

$$\mu_i = \lim_{n\to\infty} \langle \phi, \mathbf{A}'_{n,i} \rangle.$$

There exists N such that for every $n > N$ it holds

$$\sum_{i=1}^{q} |\langle \phi, \mathbf{A}'_{n,i} \rangle - \mu_i| < \epsilon/3.$$

According to Corollary 3.3 it holds, for every $n > N$:

$$\begin{aligned}
\langle\phi, \mathbf{A}_n\rangle &= \sum_i \frac{|B_{n,i}|}{|A_n|}\langle\phi, \mathbf{B}_{n,i}\rangle\\
&= \sum_{i=1}^{k} \frac{|B_{n,i}|}{|A_n|}\langle\phi, \mathbf{B}_{n,i}\rangle + \sum_{i=k+1}^{C(n)} \frac{|A_{n,i}|}{|A_n|}\langle\phi, \mathbf{B}_{n,i}\rangle + \sum_{i>C(n)} \frac{|B_{n,i}|}{|A_n|}\langle\phi, \mathbf{B}_{n,i}\rangle\\
&= \sum_{i=1}^{q} c_i\langle\phi, \mathbf{A}'_{n,i}\rangle + \sum_{i=k+1}^{C(n)} \frac{|B_{n,i}|}{|A_n|}\langle\phi, \mathbf{B}_{n,i}\rangle + \frac{|R_n|}{|A_n|}\langle\phi, \mathbf{R}_n\rangle
\end{aligned}$$

Thus we have

$$\begin{aligned}
\Big|\mathsf{sp}'\langle\phi, \mathbf{R}_n\rangle - \big(\langle\phi, \mathbf{A}_n\rangle - \sum_{i=1}^{q} c_i\mu_i\big)\Big| &\leq \sum_{i=1}^{q} |\langle\phi, \mathbf{A}'_{n,i}\rangle - \mu_i| + \sum_{i=k+1}^{C(n)} |B_{n,i}|/|A_n|\\
&\quad + \big||R_n|/|A_n| - \mathsf{sp}'\big|\\
&\leq \epsilon.
\end{aligned}$$

It follows that $\lim_{n\to\infty}\langle\phi, \mathbf{R}_n\rangle$ exists. By sorting the $C(n)$ first connected components of each $\mathbf{A}_n$ according to both sp_i and Lemma 3.49 we obtain the following expression for the limit:

$$\lim_{n\to\infty}\langle\phi, \mathbf{R}_n\rangle = \frac{1}{\mathsf{sp}'}\big(\lim_{n\to\infty}\langle\phi, \mathbf{A}_n\rangle - \sum_{i<\widehat{C}} \mathsf{sp}_i \lim_{n\to\infty}\langle\phi, \mathbf{B}_{n,i}\rangle\big).$$

□

Finally, we obtain the main results of this section.

THEOREM 3.56 (Comb structure for λ-structure sequences with infinite spectrum (local convergence)). *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be an $\mathrm{FO}^{\mathrm{local}}$-convergent sequence of finite λ-structures with component relation ϖ, support $\mathbb{N}$, and spectrum $(\mathsf{sp}_i)_{i\in\mathbb{N}}$. Let $C : \mathbb{N}\to\mathbb{N}$ be a clip of $(\mathbf{A}_n)_{n\in\mathbb{N}}$, and let $\mathbf{R}_n$ be the residue of $\mathbf{A}_n$ with respect to C.*

Then there exists, for each $n \in \mathbb{N}$, a permutation $f_n : [C(n)] \to [C(n)]$ such that, extending f_n to $\mathbb{N}$ by putting $f(i)$ to be the identity for $i > C(n)$, it holds

- $\lim_{n\to\infty}\max_{i>C(n)} |B_{n,i}|/|A_n| = 0$;
- $\mathsf{sp}' = \lim_{n\to\infty}|R_n|/|A_n|$ *exists;*
- *for every $i \in \mathbb{N}$, $(B_{n,f_n(i)})_{n\in\mathbb{N}}$ is FO-convergent;*
- *either $\mathsf{sp}' = 0$ or the sequence $(\mathbf{R}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$-convergent.*

PROOF. This lemma is a direct consequence of the previous lemmas. □

We shall now extend the Comb structure theorem to full FO-convergence. In constrast with the case of a finite ϖ-spectrum, the elementary convergence aspects will be non trivial and will require a careful choice of a clip for the sequence.

LEMMA 3.57. *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be an $\mathrm{FO}^{\mathrm{local}}$-convergent sequence of finite λ-structures with component relation ϖ, such that $\lim_{n\to\infty}|A_n| = \infty$. Let $\mathbf{B}_{n,i}$ be the connected components of $\mathbf{A}_n$. Assume that the connected components with same sp_i*

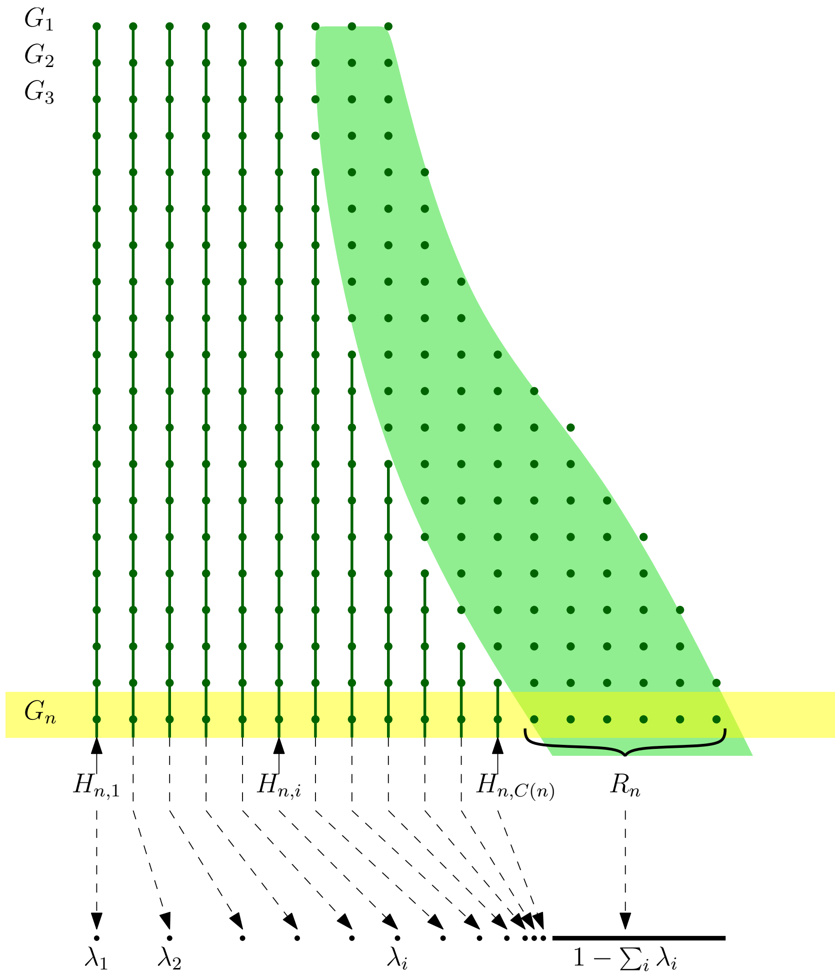

FIGURE 7. Illustration of the Comb structure theorem

have been reshuffled according to Lemma 3.49, so that $(\mathbf{B}_{n,i})_{i\in\mathbb{N}}$ *is* FO*-convergent for each* $i \in \mathbb{N}$.

For $i \in \mathbb{N}$*, let* $\widehat{\mathbf{B}_i}$ *be an elementary limit of* $(\mathbf{B}_{n,i})_{n\in\mathbb{N}}$*. Then there exists a clip* C *such that the sequence* $(\mathbf{R}_n)_{n\in\mathbb{N}}$ *of the residues is elementarily convergent. Moreover, if* $\widehat{\mathbf{R}}$ *is an elementary limit of* $(\mathbf{R}_n)_{n\in\mathbb{N}}$*, then* $\bigcup_i \widehat{\mathbf{B}_i} \cup \widehat{\mathbf{R}}$ *is an elementary limit of* $(\mathbf{A}_n)_{n\in\mathbb{N}}$.

Let $\mathbf{B}'_{n,i}$ *be either* $\mathbf{B}_{n,i}$ *if* $C(n) \geq i$ *or the empty* λ*-structure if* $C(n) < i$*. Then the family consisting in the sequences* $(\mathbf{B}'_{n,i})_{i\in\mathbb{N}}$ *(*$i \in \mathbb{N}$*) and of the sequence* $(\mathbf{R}_n)_{n\in\mathbb{N}}$ *is uniformly elementarily convergent.*

PROOF. Let $\widehat{\mathbf{A}}$ be an elementary limit of $(\mathbf{A}_n)_{n\in\mathbb{N}}$.
For $\theta \in \mathrm{FO}_1^{\mathrm{local}}$ and $m \in \mathbb{N}$ we denote by $\theta^{(m)}$ the sentence

$$\theta^{(m)}: \quad (\exists x_1 \dots \exists x_m)\left(\bigwedge_{1\leq i<j\leq m} \neg\varpi(x_i,x_j)\ \wedge\ \bigwedge_{i=1}^{m} \theta(x_i)\right).$$

According to Theorem 2.19, elementary convergence of a sequence of λ-structures with component relation ϖ can be checked by considering sentences of the form $\theta^{(k)}$ for $\theta \in \mathrm{FO}_1^{\mathrm{local}}$ and $k \in \mathbb{N}$.

Note that for every $k < k'$ and every λ-structure $\mathbf{A}$, if it holds $\mathbf{A} \models \theta^{(k')}$ then it holds $\mathbf{A} \models \theta^{(k)}$. Define

$$M(\theta) = \sup\{k \in \mathbb{N}, \quad \widehat{\mathbf{A}} \models \theta^{(k)}\}$$
$$\Omega(\theta) = \{i \in \mathbb{N}, \quad \widehat{\mathbf{B}_i} \models (\exists x)\theta(x)\}.$$

Note that obviously $|\Omega(\theta)| \leq M(\theta)$.

For $r \in \mathbb{N}$, let $\theta_1, \dots, \theta_{F(r)}$ be an enumeration of the local first-order formulas with a single free variable with quantifier rank at most r (up to logical equivalence). Define

$$A(r) = \max(r, \max_{a \leq F(r)} \max \Omega(\theta_a)).$$

Let

$$C_0(n) = \sup\left\{K, \quad (\forall n' \geq n)\ \sum_{i=1}^{K} \left| \frac{|B_{n',i}|}{|A_{n'}|} - \mathsf{sp}_i \right| \leq \sum_{i>K} \mathsf{sp}_i \right\}$$

be the standard (maximal) clip on $(\mathbf{A}_n)_{n\in\mathbb{N}}$ (see Proposition 3.2).

Let $B(r)$ be the minimum integer such that

(1) it holds $C_0(B(r)) \geq A(r)$ (note that $\lim_{n\to\infty} C_0(n) = \infty$, according to Proposition 3.2);
(2) for every $n \geq B(r), a \leq F(r)$ and every $k \leq r$ it holds $\mathbf{A}_n \models \theta_a^{(k)}$ if and only if $M(\theta_a) \geq k$ (note that this holds for sufficiently large n as $\widehat{\mathbf{A}}$ is an elementary limit of $(\mathbf{A}_n)_{n\in\mathbb{N}}$);
(3) for every $i \leq A(r)$ and $a \leq F(r)$ it holds

$$\mathbf{B}_{n,i} \models (\exists x)\theta_a(x) \quad \iff \quad \widehat{\mathbf{B}_i} \models (\exists x)\theta_a(x).$$

(note that this holds for sufficiently large n as $\widehat{\mathbf{B}_i}$ is an elementary limit of $(\mathbf{B}_{n,i})_{n\in\mathbb{N}}$ and as we consider only finitely many values of i);

we define the non-decreasing function $C : \mathbb{N} \to \mathbb{N}$ by

$$C(n) = \max\{A(r) : B(r) \leq n\}.$$

As $\lim_{r\to\infty} A(r) = \infty$ and as $C_0(B(r)) \geq A(r)$ it holds $\lim_{r\to\infty} B(r) = \infty$. Moreover, for every $r \in \mathbb{N}$ it holds $C_0(B(r)) \geq A(r)$ hence $C_0(n) \geq C(n)$. According to Proposition 3.2, it follows that the function C is a clip on $(\mathbf{A}_n)_{n\in\mathbb{N}}$.

Let $(\mathbf{R}_n)_{n\in\mathbb{N}}$ be the resiude of $(\mathbf{A}_n)_{n\in\mathbb{N}}$ with respect to the clip C, and let $\mathbf{B}'_{n,i}$ be defined as $\mathbf{B}_{n,i}$ if $i \leq C(n)$ and the empty λ-structure otherwise. Then it is immediate from the definition of the clip C that the family $\{(B'_{n,i})_{n\in\mathbb{N}} :\ i \in \mathbb{N}\}$ is uniformly elementarily convergent. Using Lemma 3.18, it is also easily checked that the residue $(\mathbf{R}_n)_{n\in\mathbb{N}}$ of $(\mathbf{A}_n)_{n\in\mathbb{N}}$ with respect to the clip C is elementarily convergent and thus, that the family $\{(B'_{n,i})_{n\in\mathbb{N}} :\ i \in \mathbb{N}\} \cup \{(\mathbf{R}_n)_{n\in\mathbb{N}}\}$ is uniformly elementarily convergent. ☐

The extension of the Comb structure theorem to FO-convergence now follows directly.

THEOREM 3.58 (Comb structure for λ-structure sequences with infinite spectrum). *Let $(\mathbf{A}_n)_{n\in\mathbb{N}}$ be an* FO*-convergent sequence of finite λ-structures with component relation ϖ and infinite spectrum $(\mathsf{sp}_i)_{i\in\mathbb{N}}$.*

Then there exists a clip $C : \mathbb{N} \to \mathbb{N}$ *with residue* $\mathbf{R}_n$ *and, for each* $n \in \mathbb{N}$*, a permutation* $f_n : [C(n)] \to [C(n)]$ *such that, extending* f_n *to* $\mathbb{N}$ *by putting* $f(i)$ *to be the identity for* $i > C(n)$*, and letting* $\mathbf{B}'_{n,i}$ *be either* $\mathbf{B}_{n,f_n(i)}$ *if* $C(n) \geq i$ *or the empty* λ*-structure if* $C(n) < i$*, it holds:*

- $\mathbf{A}_n = \mathbf{R}_n \cup \bigcup_{i \in \mathbb{N}} \mathbf{B}'_{n,i}$;
- $\lim_{n\to\infty} \max_{i > C(n)} |B'_{n,i}|/|A_n| = 0$;
- $\lim_{n\to\infty} |R_n|/|A_n|$ *exists;*
- *for every* $i \in \mathbb{N}$, $(\mathbf{B}'_{n,i})_{n\in\mathbb{N}}$ *is* FO*-convergent;*
- *either* $\lim_{n\to\infty} |R_n|/|A_n| = 0$ *and* $(\mathbf{R}_n)_{n\in\mathbb{N}}$ *is elementarily convergent, or the sequence* $(\mathbf{R}_n)_{n\in\mathbb{N}}$ *is* FO*-convergent;*
- *the family* $\{(\mathbf{B}'_{n,i})_{n\in\mathbb{N}} : i \in \mathbb{N}\} \cup \{(\mathbf{R}_n)_{n\in\mathbb{N}}\}$ *is uniformly elementarily convergent.*

This ends the (admittedly very technical and complicated) analysis of the component structure of limits. This was not developped for its own sake, but it will be all needed in the Part 3 of this paper, to construct modeling FO-limits for convergent sequences of trees with bounded height and for convergent sequences of graphs with bounded tree-depth.

CHAPTER 4

Limits of Graphs with Bounded Tree-depth

In this part, we mainly consider the signature λ, which consists in a binary relation $\sim$ (symmetric adjacency relation), a unary relation R (property of being a root), and c unary relations C_i (the coloring). Colored rooted trees with height at most h are particular λ-structures, and the class of (finite or infinite) colored rooted trees with height at most h will be denoted by $\mathcal{Y}^{(h)}$. (Here we shall be only concerned with trees that are either finite, countable, or of size continuum.)

4.1. FO_1-limits of Colored Rooted Trees with Bounded Height

In this section, we explictly define modeling FO_1-limits for FO_1-convergent sequences of colored rooted trees with bounded height and characterize modelings which are FO_1-limits for FO_1-convergent sequences of (finite) colored rooted trees with bounded height.

4.1.1. Preliminary Observations. We take some time for some preliminary observations on the logic structure of rooted colored trees with bounded height. These observations will use arguments based on Ehrenfeucht-Fraïssé games and strategy stealing. (For definitions of $\equiv^n$ and Ehrenfeucht-Fraïssé games, see Section 2.3.1.)

For a rooted colored tree $\mathbf{Y} \in \mathcal{Y}^{(h)}$ and a vertex $x \in Y$, we denote $\mathbf{Y}(x)$ the subtree of $\mathbf{Y}$ rooted at x, and by $\mathbf{Y} \setminus \mathbf{Y}(x)$ the rooted tree obtained from $\mathbf{Y}$ by removing all the vertices in $\mathbf{Y}(s)$.

The following two lemmas show that, like for isomorphism, equivalence between two colored rooted trees can be reduced to equivalence of branches.

Lemma 4.1. *Let $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}^{(h)}$, let s, s' be sons of the roots of $\mathbf{Y}$ and $\mathbf{Y}'$, respectively.*

Let $n \in \mathbb{N}$. If $\mathbf{Y}(s) \equiv^n \mathbf{Y}'(s')$ and $\mathbf{Y} \setminus \mathbf{Y}(s) \equiv^n \mathbf{Y}' \setminus \mathbf{Y}'(s')$, then $\mathbf{Y} \equiv^n \mathbf{Y}'$.

Proof. Assume $\mathbf{Y}(s) \equiv^n \mathbf{Y}'(s')$ and $\mathbf{Y} \setminus \mathbf{Y}(s) \equiv^n \mathbf{Y}' \setminus \mathbf{Y}'(s')$. In order to prove $\mathbf{Y} \equiv^n \mathbf{Y}'$ we play an n-steps Ehrenfeucht-Fraïssé-game EF_0 on $\mathbf{Y}$ and $\mathbf{Y}'$ as Duplicator. Our strategy will be based on two auxiliary n-steps Ehrenfeucht-Fraïssé-games, EF_1 and EF_2, on $\mathbf{Y}(s)$ and $\mathbf{Y}'(s')$ and on $\mathbf{Y} \setminus \mathbf{Y}(s)$ and $\mathbf{Y}' \setminus \mathbf{Y}'(s')$, respectively, against Duplicators following a winning strategy. Each time Spoiler selects a vertex in game EF_0, we play the same vertex in the game EF_1 or EF_2 (depending on the tree the vertex belongs to), then we mimic the selection of the Duplicator of this game. it is easily checked that this strategy is a winning strategy. □

Lemma 4.2. *Let $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}^{(h)}$, let s, s' be sons of the roots of $\mathbf{Y}$ and $\mathbf{Y}'$, respectively.*

Let $n \in \mathbb{N}$. If $\mathbf{Y} \equiv^{n+h} \mathbf{Y}'$ and $\mathbf{Y}(s) \equiv^n \mathbf{Y}'(s')$, then $\mathbf{Y} \setminus \mathbf{Y}(s) \equiv^n \mathbf{Y}' \setminus \mathbf{Y}'(s')$

PROOF. Assume $\mathbf{Y} \equiv^{n+h} \mathbf{Y}'$ and $\mathbf{Y}(s) \equiv^{n} \mathbf{Y}'(s')$.

We first play (as Spoiler) s in $\mathbf{Y}$ then s' in $\mathbf{Y}'$. Let t' and t be the corresponding plays of Duplicator. Then the further n steps of the game have to map vertices in $\mathbf{Y}(s)$, $\mathbf{Y}(t)$, $\mathbf{Y}\setminus(\mathbf{Y}(s)\cup\mathbf{Y}(t))$ to $\mathbf{Y}'(t')$, $\mathbf{Y}'(s')$, $\mathbf{Y}'\setminus(\mathbf{Y}'(t')\cup\mathbf{Y}'(s'))$ (and converse), for otherwise $h-2$ steps would allow Spoiler to win the game. Also, by restricting our play to one of these pairs of trees, we deduce $\mathbf{Y}(s) \equiv^{n} \mathbf{Y}'(t')$, $\mathbf{Y}(t) \equiv^{n} \mathbf{Y}'(s')$, and $\mathbf{Y} \setminus (\mathbf{Y}(s) \cup \mathbf{Y}(t)) \equiv^{n} \mathbf{Y} \setminus (\mathbf{Y}'(s') \cup \mathbf{Y}'(t'))$. As $\mathbf{Y}'(s') \equiv^{n} \mathbf{Y}(s)$ it follows

$$\mathbf{Y}(t) \equiv^{n} \mathbf{Y}'(s') \equiv^{n} \mathbf{Y}(s) \equiv^{n} \mathbf{Y}'(t').$$

Hence, according to Lemma 4.1, as $\mathbf{Y} \setminus (\mathbf{Y}(s) \cup \mathbf{Y}(t)) = (\mathbf{Y} \setminus \mathbf{Y}(s)) \setminus \mathbf{Y}(t)$ and $\mathbf{Y}'\setminus(\mathbf{Y}'(s')\cup\mathbf{Y}'(t')) = (\mathbf{Y}'\setminus\mathbf{Y}'(s'))\setminus\mathbf{Y}'(t')$, we deduce $\mathbf{Y}\setminus\mathbf{Y}(s) \equiv^{n} \mathbf{Y}'\setminus\mathbf{Y}'(s')$. ☐

Let $\lambda^{\bullet}$ denote the signature obtained from λ by adding a new unary relation S (marking a *special* vertex, which is not necessarily the root). Let $\theta_{\bullet}$ be the sentence

$$(\exists x)(S(x) \wedge (\forall y\ S(y) \rightarrow (y = x)),$$

which states that a $\lambda^{\bullet}$ contains a unique special vertex. We denote by $\mathcal{Y}_{\bullet}^{(h)}$ the class obtained by marking as special a single vertex of a colored rooted tree with height at most h. Let Unmark be the interpretation of λ-structures in $\lambda^{\bullet}$-structures consisting in forgetting S (so that Umark projects $\mathcal{Y}_{\bullet}^{(h)}$ onto $\mathcal{Y}^{(h)}$).

LEMMA 4.3. *Let $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}_{\bullet}^{(h)}$ be such that $\mathbf{Y}$ (resp. $\mathbf{Y}'$) has special vertex m (resp. m'). Assume that both m and m' have height $t > 1$ (in $\mathbf{Y}$ and $\mathbf{Y}'$, respectively). Let v (resp. v') be son of the root of $\mathbf{Y}$ (resp. $\mathbf{Y}'$) that is an ancestor of m (resp. m').*

Then for every $n \in \mathbb{N}$, if $\mathrm{Unmark}(\mathbf{Y}) \equiv^{n+h} \mathrm{Unmark}(\mathbf{Y}')$ *and* $\mathbf{Y}(v) \equiv^{n} \mathbf{Y}'(v')$, *then* $\mathbf{Y} \equiv^{n} \mathbf{Y}'$.

PROOF. Assume $\mathrm{Unmark}(\mathbf{Y}) \equiv^{n+h} \mathrm{Unmark}(\mathbf{Y}')$ and $\mathbf{Y}(v) \equiv^{n} \mathbf{Y}'(v')$. Then it holds $\mathrm{Unmark}(\mathbf{Y}(v)) \equiv^{n} \mathrm{Unmark}(\mathbf{Y}'(v'))$ thus, according to Lemma 4.2,

$$\begin{aligned}\mathbf{Y} \setminus \mathbf{Y}(v) &= \mathrm{Unmark}(\mathbf{Y}) \setminus \mathrm{Unmark}(\mathbf{Y}(v)) \\ &\equiv^{n} \mathrm{Unmark}(\mathbf{Y}') \setminus \mathrm{Unmark}(\mathbf{Y}'(v')) = \mathbf{Y}' \setminus \mathbf{Y}'(v').\end{aligned}$$

Hence, according to Lemma 4.1, it holds $\mathbf{Y} \equiv^{n} \mathbf{Y}'$ (as the marking could be considered as a coloring). ☐

The next lemma states that the properties of a colored rooted trees with a distinguished vertex v (which is not necessarily the root) can be retrieved from the properties of the subtree rooted at v, the subtree rooted at the father of v, etc. (see Fig. 1).

LEMMA 4.4. *Let $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}^{(h)}$, $v_t \in Y$ and $v_t' \in Y'$ be vertices with height t. For $1 \leq i < t$, let v_i (resp. v_i') be the ancestor of v_t (resp. of v_t') at height i.*

Then for every integer n it holds

$$(\forall 1 \leq i \leq t)\ \mathbf{Y}(v_i) \equiv^{n+h+1-i} \mathbf{Y}'(v_i') \implies (\mathbf{Y}, v_t) \equiv^{n} (\mathbf{Y}', v_t')$$

$$(\mathbf{Y}, v_t) \equiv^{n+(t-1)h} (\mathbf{Y}', v_t') \implies (\forall 1 \leq i \leq t)\ \mathbf{Y}(v_i) \equiv^{n+(t-i)h} \mathbf{Y}'(v_i')$$

PROOF. We proceed by induction over t. If $t = 1$, then the statement obviously holds. So, assume $t > 1$ and that the statement holds for $t-1$.

Let $\mathbf{Y}_{\bullet}, \mathbf{Y}_{\bullet}' \in \mathcal{Y}_{\bullet}^{(h)}$ be the marked rooted colored trees obtained from $\mathbf{Y}$ and $\mathbf{Y}'$ by marking v_t (resp. v_t') as a special vertex.

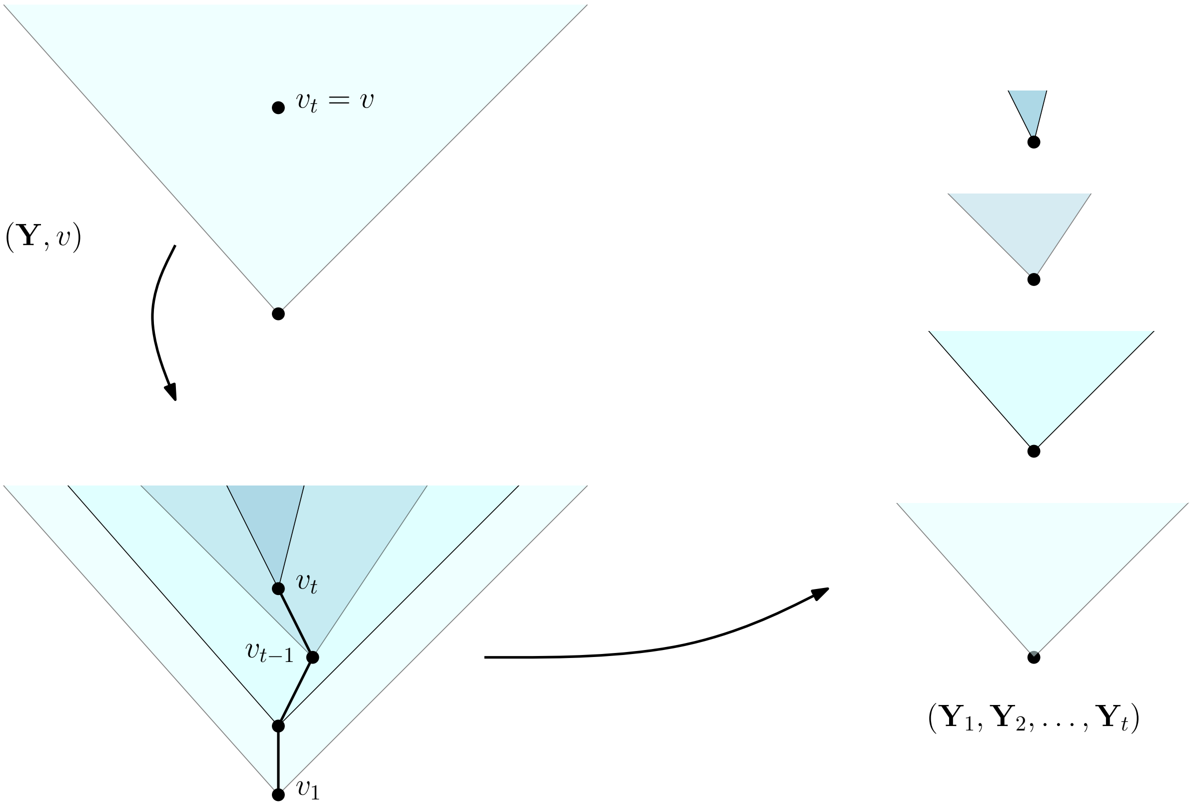

FIGURE 1. Transformation of a rooted tree with a distinguished vertex $(\mathbf{Y}, v_t)$ into a tuple of rooted trees $(\mathbf{Y}_1, \ldots, \mathbf{Y}_t)$.

Assume $(\forall 1 \leq i \leq t)\ \mathbf{Y}(v_i) \equiv^{n+h+1-i} \mathbf{Y}'(v_i')$. By induction, $(\forall 2 \leq i \leq t)\ \mathbf{Y}(v_i) \equiv^{n+(h-1)+1-(i-1)} \mathbf{Y}'(v_i')$ implies $(\mathbf{Y}(v_2), v_t) \equiv^n (\mathbf{Y}'(v_2'), v_t')$, that is $\mathbf{Y}_\bullet(v_2) \equiv^n \mathbf{Y}'_\bullet(v_2')$. As $\mathbf{Y} \equiv^{n+h} \mathbf{Y}'$, it follows from Lemma 4.3 that $\mathbf{Y}_\bullet \equiv^n \mathbf{Y}'_\bullet$, that is: $(\mathbf{Y}, v_t) \equiv^n (\mathbf{Y}', v_t')$.

Conversely, if $(\mathbf{Y}, v_t) \equiv^{n+(t-1)h} (\mathbf{Y}', v_t')$ (i.e. $\mathbf{Y}_\bullet \equiv^{n+(t-1)h} \mathbf{Y}'_\bullet$) an repeated application of Lemma 4.2 gives $\mathbf{Y}_\bullet(v_i) \equiv^{n+(t-i)h} \mathbf{Y}'_\bullet(v_i)'$ hence (by forgetting the marking) $\mathbf{Y}(v_i) \equiv^{n+(t-i)h} \mathbf{Y}'(v_i)'$. □

This lemma allows to encode the complete theory of a colored rooted tree $\mathbf{Y}$ of height at most h with special vertex v as a tuple of complete theories of colored rooted trees with height at most h.

As the height h is bounded, the classes $\mathcal{Y}^{(h)}$ can be axiomatized by finitely many axioms (hence by some single axiom $\eta_{\mathcal{Y}^{(h)}}$), it is a basic elementary class. For integer $p \geq 0$, we introduce a short notation for the Stone space associated to the Lindenbaum-Tarski algebra of formulas on $\mathcal{Y}^{(h)}$ with p free variables:

$$\mathfrak{Y}_p^{(h)} = S(\mathcal{B}(\mathrm{FO}_p(\lambda), \eta_{\mathcal{Y}^{(h)}})).$$

We shall now move from models to theories, specifically from $\mathcal{Y}_\bullet^{(h)}$ (colored rooted trees with height at most h and a special vertex) to the Stone space $\mathfrak{Y}_1^{(h)}$ and from $\mathcal{Y}^{(h)}$ (colored rooted trees with height at most h) to the Stone space $\mathfrak{Y}_0^{(h)}$.

In that direction, we first show how the notion of "property of the subtree $\mathbf{Y}(v)$ of $\mathbf{Y}$ rooted at the vertex v" translates into a relativization homomorphism $\varrho : \mathcal{B}(\mathrm{FO}_0(\lambda) \to \mathcal{B}(\mathrm{FO}_1(\lambda)$.

We consider the simple interpretation $\mathsf{I}_\bullet$ of λ-structures in $\lambda^\bullet$-structures, which maps a $\lambda^\bullet$-structure $\mathbf{Y}_\bullet$ to the λ-structure defined as follows: let $x \simeq y$ be defined as $(x \sim y) \vee (x = y)$. Then

- the domain of $\mathsf{I}_\bullet(\mathbf{Y}_\bullet)$ is defined by the formula

$$S(x_1) \vee (\forall y_1, \dots, y_h)\Big(\big(R(y_1) \wedge \bigwedge_{i=1}^{h-1} \neg S(y_i) \wedge (y_i \simeq y_{i+1})\big) \to (y_h \neq x_1)\Big);$$

- the adjacency relation $\sim$ is defined as in $\mathbf{Y}_\bullet$ (i.e. by the formula $(x_1 \sim x_2)$);
- the relation R of $\mathsf{I}_\bullet(\mathbf{Y}_\bullet)$ is defined by the formula $S(x_1)$.

Although $\mathsf{I}_\bullet$ maps general $\lambda_\bullet$-structures to λ-structures, we shall be only concerned by the specific property that $\mathsf{I}_\bullet$ maps a rooted tree $\mathbf{Y}_\bullet \in \mathcal{Y}_\bullet^{(h)}$ with special vertex v to the rooted tree $\mathrm{Unmark}(\mathbf{Y}_\bullet)(v)$.

In a sake for simplicity, for $\mathbf{Y} \in \mathcal{Y}^{(h)}$ we denote by $(\mathbf{Y}, v)$ (where $\mathbf{Y}$ is a λ-structure) the $\lambda^\bullet$-structure obtained by adding the new relation S with v being the unique special vertex.

LEMMA 4.5. *There is a Boolean algebra homomorphism*

$$\varrho : \mathcal{B}(\mathrm{FO}_0(\lambda), \eta_{\mathcal{Y}^{(h)}}) \to \mathcal{B}(\mathrm{FO}_1(\lambda), \eta_{\mathcal{Y}^{(h)}})$$

(called relativization*), such that for every sentence $\phi \in \mathrm{FO}_0(\lambda)$, every $\mathbf{Y} \in \mathcal{Y}^{(h)}$, and every $v \in Y$ it holds*

$$\mathbf{Y}(u) \models \phi \quad \iff \quad \mathbf{Y} \models \varrho(\phi)(u).$$

PROOF. The lemma follows from the property

$$\mathbf{Y}(u) \models \phi \quad \iff \quad \mathsf{I}_\bullet(\mathbf{Y}, u) \models \phi \quad \iff \quad (\mathbf{Y}, u) \models \tilde{\mathsf{I}}_\bullet(\phi).$$

The formula $\rho(\phi)$ is obtained from the sentence $\tilde{\mathsf{I}}_\bullet(\phi)$ by replacing each occurence of $S(y)$ by $y = x_1$. □

Using relativization and Lemma 4.4, we can translate the transformation shown on Figure 1 to a encoding of elements of $\mathfrak{Y}_1^{(h)}$ into tuples of elements $\mathfrak{Y}_0^{(h)}$. Intuitively, a element $T \in \mathfrak{Y}_1^{(h)}$ defines the properties of a colored rooted tree $\mathbf{Y}$ with special vertex x_1, and the relativization ρ allows us to extract from T the tuple of the complete theories of the subtrees of $\mathbf{Y}$ rooted at x_1, the father of x_1, etc. Moreover, the meaning of Lemma 4.4 is that what we obtain only depends on the complete theory of $(\mathbf{Y}, x_1)$, that is only on T.

DEFINITION 4.6. For $1 \le i \le h$, let $\eta_i \in \mathrm{FO}_1(\lambda)$ be the formula stating that the height of x_1 is i.

We define the mapping $\mathrm{Encode} : \mathfrak{Y}_1^{(h)} \to \biguplus_{k=1}^{h} (\mathfrak{Y}_0^{(h)})^k$ as follows:

For $T \in \mathfrak{Y}_1^{(h)}$, let k be the (unique) integer such that $\eta_k \in T$. Then $\mathrm{Encode}(T)$ is the k-tuple $(T_0, \dots, T_{k-1})$, where

- T_{k-1} is the set of sentences $\theta \in \mathrm{FO}_0(\lambda)$ such that $\rho(\theta) \in T$ (intuitively, the complete theory of the subtree rooted at x_1);
- T_{k-2} is the set of sentences $\theta \in \mathrm{FO}_0(\lambda)$ such that

$$\big((\exists y_1)(\eta_{k-1}(y_1) \wedge y_1 \sim x_1 \wedge \rho(\theta)(y)\big) \in T$$

  (intuitively, the complete theory of the subtree rooted at father of x_1);

- T_{k-1-i} is the set of sentences $\theta \in \mathrm{FO}_0(\lambda)$ such that
$$\big((\exists y_1 \dots y_i)(\bigwedge_{j=1}^{i} \eta_{k-j}(y_j) \wedge \bigwedge_{j=1}^{i-1}(y_j \sim y_{j+1} \wedge y_1 \sim x_1 \wedge \rho(\theta)(y_i)\big) \in T$$
(intuitively, the complete theory of the subtree rooted at the ancestor of x_1 which has height $k-i$);
- $T_0 = T \cap \in \mathrm{FO}_0(\lambda)$ (intuitively, the complete theory of the whole rooted tree).

LEMMA 4.7. Encode *is a homeomorphism of* $\mathfrak{Y}_1^{(h)}$ *and* $\mathrm{Encode}(\mathfrak{Y}_1^{(h)})$*, which is a closed subspace of* $\biguplus_{k=1}^{h}(\mathfrak{Y}_0^{(h)})^k$*.*

PROOF. This lemma is a direct consequence of Lemma 4.4. $\square$

4.1.2. The Universal Relational Sample Space $\mathbb{Y}_h$. The aim of this section is to construct a rooted colored forest on a standard Borel space $\mathbb{Y}_h$ that is FO_1-*universal*, in the sense that every FO_1-convergent sequence of colored rooted trees will have a modeling FO_1-limit obtained by assigning an adapted probability measure to one of the connected components of $\mathbb{Y}_h$.

DEFINITION 4.8. For theories $T, T' \in \mathfrak{Y}_0^{(h)}$, we define $w(T,T') \geq k$ if and only if there exists a model $\mathbf{Y}$ of T, such that the root of $\mathbf{Y}$ has k (distinct) sons $v_1,\dots,v_k$ with $\mathrm{Th}(\mathbf{Y}(v_i)) = T'$.

LEMMA 4.9. *For* $k \in \mathbb{N}$ *and* $\phi \in \mathrm{FO}_0$*, let* $\zeta_k(\phi)$ *be the sentence* $(\exists^{\geq k} y)\ \rho(\phi)(y)$*. Then* $w(T,T') \geq k$ *if and only if* $\zeta_k(\phi) \in T$ *holds for every* $\phi \in T'$*.*

PROOF. If $w(T,T') \geq k$, then $\zeta_k(\phi) \in T$ holds for every $\phi \in T'$, hence we only have to prove the opposite direction. Assume that $\zeta_k(\phi) \in T$ holds for every $\phi \in T'$, but that there is $\phi_0 \in T'$ such that $\zeta_{k+1}(\phi_0) \notin T$. Let Y be a model of T, and let $v_1,\dots,v_k$ be the sons of the root of Y such that $Y(v_i) \models \phi_0$. For every $r \in \mathbb{N}$, $r \geq \mathrm{qrank}(\phi_0)$, let ψ_r be the conjunction of the sentences in T' with quantifier rank r. Obviously, $\psi_r \in T'$. Moreover, as $\zeta_k(\psi_r) \in T$, it holds $Y \models \zeta_k(\psi_r)$. As $\psi_r \vdash \phi_0$, it follows that for every $1 \leq i \leq k$ it holds $Y(v_i) \models \psi_r$ (only possible choices). As this holds for every r, we infer that for every $1 \leq i \leq k$, $Y(v_i)$ is a model of T' hence $w(T,T') \geq k$. Now assume that for every $k \in \mathbb{N}$ and every $\phi \in T'$ it holds $\zeta_k(\phi) \in T$. Let Y be a model of T, let Y' be a model of T', and let $\tilde{Y}$ be obtained from Y by adding (at the root of Y) a son u with subtree $\tilde{Y}(u)$ isomorphic to Y'. By an easy application of an Ehrenfeucht-Fraïssé game, we get that $\tilde{Y}$ is elementarily equivalent to Y, hence a model of T. Thus $w(T,T') \geq k$. $\square$

Let $\overline{\mathbb{N}}$ be the one point compactification of $\mathbb{N}$, that is $\overline{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$ with open sets generated by complements of finite sets.

LEMMA 4.10. *The function* $w : \mathfrak{Y}_0^{(h)} \times \mathfrak{Y}_0^{(h)} \to \overline{\mathbb{N}}$ *is upper semicontinuous (with respect to product topology of Stone space* $\mathfrak{Y}_0^{(h)}$*).*

PROOF. For $r \in \mathbb{N}$ define the function $w_r : \mathfrak{Y}_0^{(h)} \times \mathfrak{Y}_0^{(h)} \to \overline{\mathbb{N}}$ by:
$$w_r(T,T') = \sup\{k \in \mathbb{N} :\ \forall \psi \in T'\ (\mathrm{qrank}(\psi) \leq r) \Rightarrow \zeta_k(\psi) \in T\}.$$
It follows from Lemma 4.9 that it holds
$$w(T,T') = \inf_{r \in \mathbb{N}} w_r(T,T').$$

Hence, in order to prove that the function w is upper semicontinuous, it is sufficient to prove that functions w_r are continuous. Let $(T_0, T_0') \in \mathfrak{Y}_0^{(h)} \times \mathfrak{Y}_0^{(h)}$.
– Assume $w_r(T_0, T_0') = k$.
If $\mathrm{dist}(T', T_0') < 2^{-r}$ and $\mathrm{dist}(T, T_0) < 2^{-\max\{\mathrm{qrank}(\zeta_{k+1}(\psi)\colon \mathrm{qrank}(\psi) \leq r\}}$, then it holds $w_r(T, T') = w_r(T_0, T_0')$;
– Assume $w_r(T_0, T_0') = \infty$, and let $k \in \mathbb{N}$.
If $\mathrm{dist}(T', T_0') < 2^{-r}$ and $\mathrm{dist}(T, T_0) < 2^{-\max\{\mathrm{qrank}(\zeta_{k+1}(\psi)\colon \mathrm{qrank}(\psi) \leq r\}}$, then it holds $w_r(T, T') > k$. □

For $z = (z_1, \dots, z_a) \in \overline{\mathbb{N}}^a$ define the subset F_z of $(\mathfrak{Y}_0^{(h)})^{a+1}$ by

$$F_z = \{(T_0, \dots, T_a) : w(T_{i-1}, T_i) = z_i\}.$$

For $t \in \overline{\mathbb{N}}$, define

$$X_t = \begin{cases} \{1, \dots, t\}, & \text{if } t \in \mathbb{N}, \\ [0,1], & \text{if } t = \infty. \end{cases}$$

For $z = (z_1, \dots, z_a) \in \overline{\mathbb{N}}^a$, define $X_z = \prod_{i=1}^a X_{z_i}$. Let

$$V_h = \mathfrak{Y}_0^{(h)} \uplus \biguplus_{a=1}^{h-1} \biguplus_{z \in \overline{\mathbb{N}}^a} (F_z \times X_z).$$

DEFINITION 4.11. The *universal forest* $\mathbb{Y}_h$ has vertex set V_h. The roots of $\mathbb{Y}_h$ are the elements in $\mathfrak{Y}_0^{(h)}$. The edges of $\mathbb{Y}_h$ are the pairs of the form

$$\{((T_0, T_1, \dots, T_a), (\alpha_1, \dots, \alpha_a)), ((T_0, T_1, \dots, T_{a+1}), (\alpha_1, \dots, \alpha_{a+1}))\}$$

where $T_i \in \mathfrak{Y}_0^{(h)}$, $\alpha_i \in [0,1]$ and $a \in \{0, \dots, h-1\}$.

Moreover, the vertex set V_h inherits the topological structure of $\biguplus_{i=1}^h (\mathfrak{Y}_0^{(h)})^i \times [0,1]^{i-1}$, which defines a σ-algebra Σ_h on V_h (as the trace on V_h of the Borel σ-algebra of $\biguplus_{i=1}^h (\mathfrak{Y}_0^{(h)})^i \times [0,1]^{i-1}$).

REMARK 4.12. Let T_0 be the complete complete theory of a colored rooted tree with height at most h. Then, by construction, T_0 is the complete theory of the connected component of $\mathbb{Y}_h$ rooted at T_0. In particular, no two connected components of $\mathbb{Y}_h$ are elementarily equivalent.

The remaining of this section will be devoted to the proof of Theorem 4.14, which states that $\mathbb{Y}_h$ is a relational sample space. In order to prove this result, we shall need a preliminary lemma, which expresses that the property of a tuple of vertices in a colored rooted tree with bounded height is completely determined by the individual properties of the vertices in the tuple and the heights of the lowest common ancestors of every pair of vertices in the tuples.

LEMMA 4.13. *Fix rooted trees $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}^{(h)}$. Let $u_1, \dots, u_p$ be p vertices of $\mathbf{Y}$, let $u_1', \dots, u_p'$ be p vertices of $\mathbf{Y}'$, and let $n \in \mathbb{N}$.*

Assume that for every $1 \leq i \leq p$ it holds $(\mathbf{Y}, u_i) \equiv^{n+h} (\mathbf{Y}', u_i')$ and that for every $1 \leq i, j \leq p$ the height of $u_i \wedge u_j$ in $\mathbf{Y}$ is the same as the height of $u_i' \wedge u_j'$ in $\mathbf{Y}'$ (where $u \wedge v$ denotes the lowest common ancestor of u and v).

Then $(\mathbf{Y}, u_i, \dots, u_p) \equiv^n (\mathbf{Y}', u_1', \dots, u_p')$.

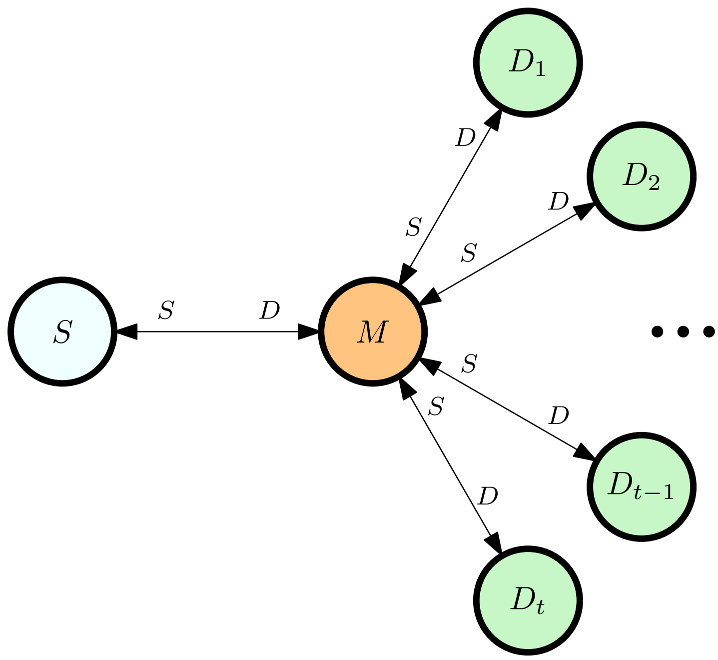

FIGURE 2. A winning strategy for $\mathrm{EF}((Y, u_1, \dots, u_p), (Y', u'_1, \dots, u'_p), n)$ using p auxiliary games $\mathrm{EF}((Y, u_i), (Y', u'_i), n + h)$.

PROOF. In the proof we consider $p+1$ simultaneous Ehrenfeucht-Fraïssé games (see Fig. 2).

Consider an n-step Ehrenfeucht-Fraïssé $\mathrm{EF}((\mathbf{Y}, u_1, \dots, u_p), (\mathbf{Y}', u'_1, \dots, u'_p), n)$ on $(\mathbf{Y}, u_1, \dots, u_p)$ and $(\mathbf{Y}', u'_1, \dots, u'_p)$. We build a strategy for Duplicator by considering p auxiliary Ehrenfeucht-Fraïssé games $\mathrm{EF}((\mathbf{Y}, u_i), (\mathbf{Y}', u'_i), n + h)$ on $(\mathbf{Y}, u_i)$ and $(\mathbf{Y}', u'_i)$ (for $1 \le i \le p$) where we play the role of Spoiler against Duplicators having a winning strategy for $n + h$ steps games.

For every vertex $v \in Y$ (resp. $v' \in Y'$) let $p(v)$ (resp. $p'(v)$) be the maximum ancestor of v (in the sense of the furthest from the root) such that $p(v) \le u_i$ (resp. $p'(v) \le u'_i$) for some $1 \le i \le p$. We partition Y and Y' as follows: for every vertex $v \in Y$ (resp. $v' \in Y'$) we put $v \in V_i$ (resp. $v' \in V'_i$) if i it the minimum integer such that $p(v) \le u_i$ (resp. such that $p'(v) \le u'_i$), see Fig 3.

Note that each V_i (resp. V'_i induces a connected subgraph of $\mathbf{Y}$ (resp. of $\mathbf{Y}'$).

Assume that at round $j \le n$, Spoiler plays a vertex $v \in (\mathbf{Y}, u_1, \dots, u_p)$ (resp. a vertex $v' \in (\mathbf{Y}', u'_1, \dots, u'_p)$).

If $v \in V_i$ (resp. $v' \in V'_i$) for some $1 \le i \le p$ then we play v (resp. v') on $(\mathbf{Y}, u_i)$ (resp. $(\mathbf{Y}', u'_i)$). We play Duplicator on $(\mathbf{Y}', u'_1, \dots, u'_p)$ (resp. on $(\mathbf{Y}, u_1, \dots, u_p)$) with the same move as our Duplicator opponent did on $(\mathbf{Y}', u_i)$ (resp. on $(\mathbf{Y}, u_i)$). If all the Duplicators' are not form a coherent then it is easily checked that h additional moves (at most) are sufficient for at least one of the Spoilers to win one of the p games, contradicting the hypothesis of p winning strategies for Duplicators. It follows that $(\mathbf{Y}, u_i, \dots, u_p) \equiv^n (\mathbf{Y}', u'_1, \dots, u'_p)$. □

THEOREM 4.14. *The rooted colored forest $\mathbb{Y}_h$ (equipped with the σ-algebra Σ_h) is a relational sample space.*

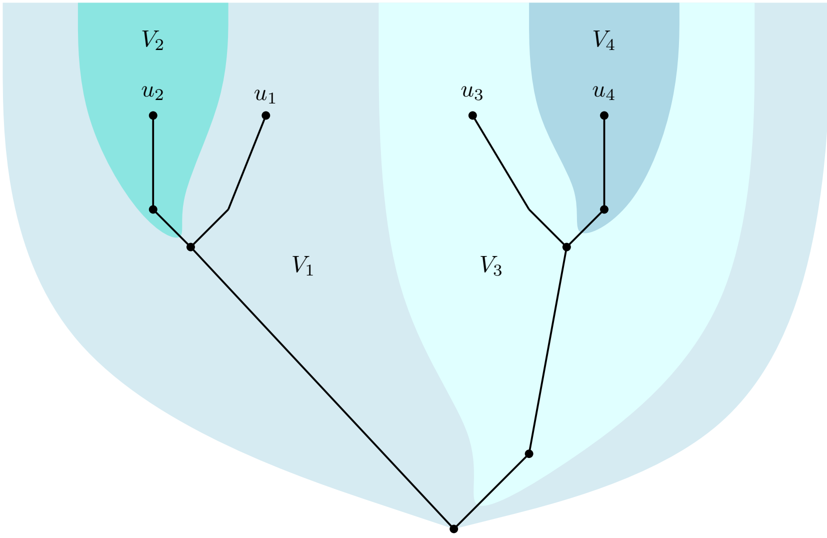

FIGURE 3. The partition (V_1, V_2, V_3, V_4) of Y induced by (u_1, u_2, u_3, u_4).

PROOF. It suffices to prove that for every $p \in \mathbb{N}$ and every $\varphi \in \mathrm{FO}_p$ the set

$$\Omega_\varphi(\mathbb{Y}_h) = \{(v_1, \dots, v_p) \in V_h^p : \ \mathbb{Y}_h \models \varphi(v_1, \dots, v_p)\}$$

is measurable.

Let $\varphi \in \mathrm{FO}_p$ and let $n = \mathrm{qrank}(\varphi)$.

We partition V_h into equivalence classes modulo $\equiv^{n+h}$, which we denote $C_1, \dots, C_N$.

Let $i_1, \dots, i_p \in [N]$ and, for $1 \le j \le p$, let v_j and v'_j belong to C_{i_j}.

According to Lemma 4.13, if the heights of the lowest common ancestors of the pairs in $(v_1, \dots, v_p)$ coincide with the heights of the lowest common ancestors of the pairs in $(v'_1, \dots, v'_p)$ then it holds

$$(\mathbb{Y}_h, v_1, \dots, v_p) \equiv^n (\mathbb{Y}_h, v'_1, \dots, v'_p)$$

thus $(v_1, \dots, v_p) \in \Omega_\varphi(\mathbb{Y}_h)$ if and only if $(v'_1, \dots, v'_p) \in \Omega_\varphi(\mathbb{Y}_h)$.

It follows from Lemma 4.7 (and the definition of $\mathbb{V}_h$ and Σ_h) that each C_j is measurable. According to Lemma 4.7 and the encoding of the vertices of V_h, the conditions on the heights of lowest common ancestors rewrite as equalities and inequalities of coordinates. It follows that $\Omega_\varphi(\mathbb{Y}_h)$ is measurable. $\square$

4.1.3. Modeling FO_1-limits of Colored Rooted Trees with Bounded Height. Let $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ be an FO_1-convergent of colored rooted trees with height at most h, and let $\widetilde{\mathbf{Y}}$ be the connected component of $\mathbb{Y}_h$ that is an elementary limit of $(\mathbf{Y}_n)_{n\in\mathbb{N}}$. According to Lemma 3.3, $\widetilde{\mathbf{Y}}$ is a relational sample space. We have to transfer the measure μ we obtained in Theorem 2.7 on $S(\mathcal{B}(\mathrm{FO}_1))$ to $\widetilde{\mathbf{Y}}$.

DEFINITION 4.15. Let μ be a measure on $\mathfrak{Y}_1^{(h)}$. We define ν on $\mathbb{Y}_h$ as follows: let $\widetilde{\mu} = \mathrm{Encode}_*(\mu)$ be the pushforward of μ by Encode (see page 11). For $t \in \overline{\mathbb{N}}$ we equip X_t with uniform discrete probability measure if $t < \infty$ and the Haar probability measure if $t = \infty$. For $z \in \overline{\mathbb{N}}^a$, X_z is equipped with the corresponding product measure, which we denote by λ_z (not to be confused with signature λ).

We define the measure ν as follows: let A be a measurable subset of V_h, let $A_0 = A \cap \mathfrak{Y}_0^{(h)}$, and let $A_z = A \cap (F_z \times X_z)$. Then

$$\nu(A) = \widetilde{\mu}(A_0) + \sum_{a=1}^{h-1} \sum_{z \in \overline{\mathbb{N}}^a} (\widetilde{\mu} \otimes \lambda_z)(A_z).$$

(Notice that the sets A_z are measurable as $F_z \times X_z$ is measurable for every z.)

LEMMA 4.16. *The measure μ is the push-forward of ν by the projection $P : \mathbb{Y}_h \to \mathfrak{Y}_1^{(h)}$ defined by*

$$P((T_0, T_1, \alpha_1, \dots, T_a, \alpha_a)) = \mathrm{Encode}^{-1}(T_0, \dots, T_a),$$

that is: $\mu = P_*(\nu)$.

PROOF. First notice that P is continuous, as Encode is a homeomorphism (by Lemma 4.7). Let B be a measurable set of $\mathfrak{Y}_1^{(h)}$. Let $A = P^{-1}(B)$. Then $A \cap (F_z \times X_z) = (\mathrm{Encode}(B) \cap F_z) \times X_z$ hence

$$(\widetilde{\mu} \otimes \lambda_z)(A \cap (F_z \times X_z)) = \nu(\mathrm{Encode}(B) \cap F_z)\lambda_z(X_z) = \widetilde{\mu}(\mathrm{Encode}(B) \cap F_z).$$

It follows that

$$\begin{aligned}
P_*(\nu)(B) &= \nu(A) \\
&= \widetilde{\mu}(A \cap \mathfrak{Y}_0^{(h)}) + \sum_{a=1}^{h-1} \sum_{z \in \overline{\mathbb{N}}^a} (\widetilde{\mu} \otimes \lambda_z)(A \cap (F_z \times X_z)) \\
&= \widetilde{\mu}(\mathrm{Encode}(B) \cap \mathfrak{Y}_0^{(h)}) + \sum_{a=1}^{h-1} \sum_{z \in \overline{\mathbb{N}}^a} \widetilde{\mu}(\mathrm{Encode}(B) \cap F_z) \\
&= \widetilde{\mu}\big(\mathrm{Encode}(B) \cap (\mathfrak{Y}_0^{(h)} \uplus \biguplus_{a=1}^{h-1} \biguplus_{z \in \overline{\mathbb{N}}^a} F_z)\big) \\
&= \widetilde{\mu} \circ \mathrm{Encode}(B) \\
&= \mu(B).
\end{aligned}$$

(as z ranges over a countable set and as all the F_z are measurable). Hence $\mu = P_*(\nu)$. ∎

LEMMA 4.17. *Let μ be a pure measure on $\mathfrak{Y}_1^{(h)}$ and let T_0 the complete theory of μ (see Definition 2.9). Let ν be the measure defined from μ by Definition 4.15. Let $\widetilde{\mathbf{Y}}$ be the connected component of $\mathbb{Y}_h$ containing the support of ν. Let $\nu_{\widetilde{\mathbf{Y}}}$ be the restriction of ν to $\widetilde{\mathbf{Y}}$.*

Then $\widetilde{\mathbf{Y}}$, equipped with the probability measure $\nu_{\widetilde{\mathbf{Y}}}$ is a modeling such that for every $\varphi \in \mathrm{FO}_1$ it holds

$$\langle \varphi, \widetilde{\mathbf{Y}} \rangle = \mu(K(\varphi)).$$

Let $X \subset \mathfrak{Y}_1^{(h)}$ be the set of all $T \in \mathfrak{Y}_1^{(h)}$ such that x_1 is not the root (i.e. $X = \{T \in \mathfrak{Y}_1^{(h)} : R(x_1) \notin T\}$). Let $f : X \to \mathfrak{Y}_0^{(h)}$ be the second projection of Encode *(if* $\mathrm{Encode}(T) = (T_0, \dots, T_i)$ *then $f(T) = T_1$). Let $\kappa = f_*(\mu)$ be the pushforward of μ by f. Intuitively, for $T \in \mathfrak{Y}_0^{(h)}$, $\kappa(\{T\})$ is the global measure of all the subtrees with complete theory T that are rooted at a son of the root.*

Let $r_{\widetilde{\mathbf{Y}}}$ be the root of $\widetilde{\mathbf{Y}}$. Then it holds

$$\sup_{v\sim r_{\widetilde{\mathbf{Y}}}} \nu_{\widetilde{\mathbf{Y}}}(\widetilde{\mathbf{Y}}(v)) = \sup_{T\in X} \frac{\kappa(\{T\})}{w(T_0,T)}.$$

PROOF. As μ is pure, the complete theory of μ is the theory T_0 to which every point of the support of μ projects. Hence the support of μ defines a unique connected component $\widetilde{\mathbf{Y}}$ of $\mathbb{Y}_h$. That for every $\varphi \in \mathrm{FO}_1$ it holds

$$\langle \varphi, \widetilde{\mathbf{Y}}\rangle = \mu(K(\varphi))$$

is a direct consequence of Lemma 4.16.

The second equation is a direct consequence of the construction of $\nu_{\widetilde{\mathbf{Y}}}$. □

THEOREM 4.18. *Let $\mathbf{Y}_n$ be an FO_1-convergent sequence of colored rooted trees with height at most h, and let μ be the limit measure of $\mu_{\mathbf{Y}_n}$ on $\mathfrak{Y}_1^{(h)}$. Let ν be the measure defined from μ by Definition 4.15. Let $\widetilde{\mathbf{Y}}$ be the connected component of $\mathbb{Y}_h$ containing the support of ν. Let $\nu_{\widetilde{\mathbf{Y}}}$ be the restriction of ν to $\widetilde{\mathbf{Y}}$.*

Then $\widetilde{\mathbf{Y}}$, equipped with the probability measure $\nu_{\widetilde{\mathbf{Y}}}$, is a modeling FO_1-limit of $(\mathbf{Y}_n)_{n\in\mathbb{N}}$.

Moreover, it holds

$$\sup_{v\sim r_{\widetilde{\mathbf{Y}}}} \nu_{\widetilde{\mathbf{Y}}}(\widetilde{\mathbf{Y}}(v)) \leq \liminf_{n\to\infty} \max_{v\sim r_{\mathbf{Y}_n}} \frac{|\mathbf{Y}_n(v)|}{|\mathbf{Y}_n|}.$$

PROOF. As $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is elementarily convergent, the complete theory of the elementary limit of this sequence is the theory T_0 to which every point of the support of μ projects, hence μ is pure. According to Lemma 4.17, $\widetilde{\mathbf{Y}}$ is an FO_1-modeling limit of $(\mathbf{Y}_n)_{n\in\mathbb{N}}$.

Let κ be defined as in Lemma 4.17. If κ is atomless, then $\sup_{v\sim r_{\widetilde{\mathbf{Y}}}} \nu_{\widetilde{\mathbf{Y}}}(\widetilde{\mathbf{Y}}(v)) = 0$ hence the inequation holds.

Let T be such that $\kappa(\{T\}) > 0$. For every $\epsilon > 0$ there exists $\theta_\epsilon \in T$ such that

$$\kappa(\{T\}) \leq \kappa(\{T' :\ T' \ni \theta_\epsilon\}) \leq \kappa(\{T\}) + \epsilon.$$

Moreover, it follows from Lemma 4.9 that

$$w(T_0,T) = \lim_{\epsilon\to 0} \sum \{w(T_0,T') :\ T' \ni \theta_\epsilon\}.$$

Then, as $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is elementarily convergent to a rooted tree with theory T_0 it holds

$$\begin{aligned} w(T_0,T) &= \lim_{\epsilon\to 0} \lim_{n\to\infty} |\{v \in Y_n :\ v \sim r_{\mathbf{Y}_n} \text{ and } \mathbf{Y}_n(v) \models \theta_\epsilon\}| \\ &= \lim_{\epsilon\to 0} \lim_{n\to\infty} |\{v \in Y_n :\ \mathbf{Y}_n \models (v \sim r_{\mathbf{Y}_n}) \wedge \rho(\theta_\epsilon)(v)\}|. \end{aligned}$$

As $\lim_{n\to\infty} |\{v \in Y_n :\ \mathbf{Y}_n \models (v \sim r_{\mathbf{Y}_n}) \wedge \rho(\theta_\epsilon)(v)\}|$ is non-increasing when $\epsilon \to 0$, and is a non-negative integer or ∞, there exists ϵ_0 such that for every $0 < \epsilon < \epsilon_0$ it holds

$$w(T_0,T) = \lim_{n\to\infty} |\{v \in Y_n :\ \mathbf{Y}_n \models (v \sim r_{\mathbf{Y}_n}) \wedge \rho(\theta_\epsilon)(v)\}|.$$

For $\epsilon > 0$, let ϕ_ϵ be the formula stating that the subtree rooted at a son of the root that contains x_1 satisfies θ_ϵ. Then it holds

$$\begin{aligned}\kappa(\{T' : T' \ni \theta_\epsilon\}) &= \mu(K(\phi_\epsilon))\\ &= \lim_{n\to\infty} \langle \phi_\epsilon, \mathbf{Y}_n\rangle\\ &= \lim_{n\to\infty} \frac{\sum\{|\mathbf{Y}_n(v)| :\ \mathbf{Y}_n \models (v \sim r_{\mathbf{Y}_n}) \wedge \rho(\theta_\epsilon)(v)\}}{|\mathbf{Y}_n|}\end{aligned}$$

Hence, for every $0 < \epsilon < \epsilon_0$ it holds

$$\begin{aligned}\frac{\kappa(\{T\})}{w(T_0,T)} &\leq \lim_{n\to\infty} \frac{\epsilon + \sum\left\{\frac{|\mathbf{Y}_n(v)|}{|\mathbf{Y}_n|} :\ \mathbf{Y}_n \models (v \sim r_{\mathbf{Y}_n}) \wedge \rho(\theta_\epsilon)(v)\right\}}{|\{v \in Y_n :\ \mathbf{Y}_n \models (v \sim r_{\mathbf{Y}_n}) \wedge \rho(\theta_\epsilon)(v)\}|}\\ &\leq \epsilon + \liminf_{n\to\infty} \max\left\{\frac{|\mathbf{Y}_n(v)|}{|\mathbf{Y}_n|} :\ \mathbf{Y}_n \models (v \sim r_{\mathbf{Y}_n}) \wedge \rho(\theta_\epsilon)(v)\right\}\\ &\leq \epsilon + \liminf_{n\to\infty} \max_{v\sim r_{\mathbf{Y}_n}} \frac{|\mathbf{Y}_n(v)|}{|\mathbf{Y}_n|}.\end{aligned}$$

Hence

$$\frac{\kappa(\{T\})}{w(T_0,T)} \leq \liminf_{n\to\infty} \max_{v\sim r_{\mathbf{Y}_n}} \frac{|\mathbf{Y}_n(v)|}{|\mathbf{Y}_n|}.$$

$\square$

4.1.4. Inverse Theorems for FO_1-limits of Colored Rooted Trees with Bounded Height. We characterize here the measures μ on $S(\mathcal{B}(\mathrm{FO}_1))$, which are weak limits of measures μ_{Y_n} for some FO_1-convergent sequence $(Y_n)_{n\in\mathbb{N}}$ of colored rooted trees with height at most h.

FACT 4.19. If $(Y_n)_{n\in\mathbb{N}}$ is an FO_1-convergent sequence of colored rooted trees with height at most h, then μ is pure and its complete theory is the limit in $S(\mathbf{B}(\mathrm{FO}_0))$ of the complete theories of the rooted trees Y_n.

We need the following combinatorial condition for a probability measure on a Stone space.

DEFINITION 4.20. A probability measure μ on $S(\mathcal{B}(\mathrm{FO}_p(\lambda)))$ ($p \geq 1$) or $S(\mathcal{B}(\mathrm{FO}(\lambda)))$ satisfies the *Finitary Mass Transport Principle* (FMTP) if for every $\phi, \psi \in \mathrm{FO}_1(\lambda)$ and every integers a, b such that

$$\begin{cases}\phi \vdash (\exists^{\geq a} y)\,(x_1 \sim y) \wedge \psi(y)\\ \psi \vdash (\exists^{\leq b} y)\,(x_1 \sim y) \wedge \psi(y)\end{cases}$$

it holds

$$a\,\mu(K(\phi)) \leq b\,\mu(K(\psi)).$$

Similarly, a modeling $\mathbf{L}$ satisfies the FMTP if, for every ϕ, ψ, a, b as above it holds (see Fig. 4):

$$a\langle \phi, \mathbf{L}\rangle \leq b\langle \psi, \mathbf{L}\rangle.$$

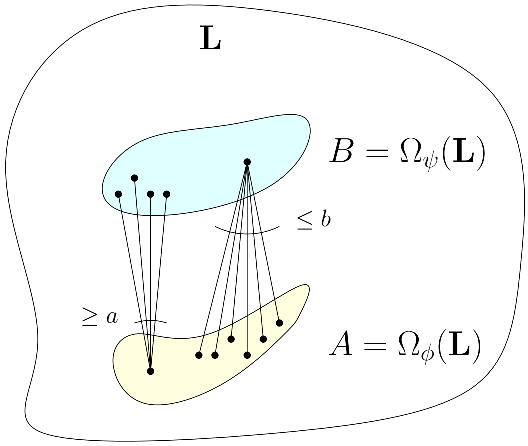

$\Longrightarrow \quad a\,\nu_{\mathbf{L}}(A) \leq b\,\nu_{\mathbf{L}}(B)$

FIGURE 4. A modeling $\mathbf{L}$ satisfies the FMTP if, for every first-order definable subsets A, B of L and every integers a, b with the property that every element in A has at least b neighbours in B and every element in B has at most b neighbours in A, it holds $a\,\nu_{\mathbf{L}}(A) \leq b\,\nu_{\mathbf{L}}(B)$.

FACT 4.21. Every finite structure $\mathbf{A}$ satisfies the FMTP and, consequently, the measures $\mu_{\mathbf{A}}$ associated to $\mathbf{A}$ on $S(\mathcal{B}(\mathrm{FO}_p))$ ($p \geq 1$) and $S(\mathcal{B}(\mathrm{FO}))$ satisfy the FMTP.

Let $r \in \mathbb{N}$. we denote by $\mathrm{FO}_1^{(r)}$ the fragment of FO_1 with formulas having quantifier-rank at most r. Note that $\mathcal{B}(\mathrm{FO}_1^{(r)})$ is a finite Boolean algebra, hence $S(\mathcal{B}(\mathrm{FO}_1^{(r)}))$ is a finite space.

The following approximation lemma lies in the centre of our inverse argument.

LEMMA 4.22. *Let μ be a pure measure on $S(\mathcal{B}(\mathrm{FO}_1(\lambda)))$ with support in $\mathfrak{Y}_1^{(h)}$ that satisfies the FMTP. Then, for every integer $r \geq 1$ there exist integer $C = C(\lambda, r)$ such that for every $N \in \mathbb{N}$ there is a colored rooted tree Y_N with the following properties:*

(1) $N \leq |Y_N| \leq N + C$;

(2) for every $\varphi \in \mathrm{FO}_1$ with quantifier rank at most r it holds

$$\left|\langle \varphi, Y_N\rangle - \mu(K(\varphi))\right| \leq C/N.$$

(3) the trees Y_N (with root r_N) are balanced in the following sense: for every modeling $\mathbf{L}$ (with root $r_{\mathbf{L}}$) such that $\langle \phi, \mathbf{L}\rangle = \mu(K(\phi))$ holds for every $\phi \in \mathrm{FO}_1$, we have

$$\max_{v \sim r_N} \frac{|Y_N(v)|}{|Y_N|} \leq \max\Bigl(\frac{1}{r+h}, \sup_{v \sim r_{\mathbf{L}}} \nu_{\mathbf{L}}(\mathbf{L}(v))\Bigr) + C/N.$$

PROOF. Remark that it is sufficient to prove that there exists C such that for every $N \geq C$ the statement holds (then the initial statement obviously holds with constant $2C$ instead of C).

For integers k, r and a sentence $\phi \in \mathrm{FO}_0(\lambda)$, let $\zeta_k(\phi)$ be the sentence

$$(\exists^{\geq k} y)\ \varrho(\phi)(y),$$

and let

$$c(s, k) = k + 1 + \max\{\mathrm{qrank}(\zeta_k(\phi)) :\ \phi \in \mathrm{FO}_0(\lambda) \text{ and } \mathrm{qrank}(\phi) \leq s\}.$$

For formulas ϕ, ψ we define

$$w'(\phi,\psi) = \begin{cases} 0 & \text{if } \phi \vdash \nexists y \ \varrho(\psi)(y) \\ k & \text{if } 0 < k < r+h, \ \phi \vdash \zeta_k(\psi), \text{ and } \phi \nvdash \zeta_{k+1}(\psi) \\ r+h & \text{otherwise.} \end{cases}$$

Let $T, T' \in \mathfrak{Y}_0^{(h)}$ be complete theories of rooted trees, let a, b are integers such that $a \geq c(b, r+h)$, let $\phi = \bigwedge(T \cap \mathrm{FO}_0^{(a)})$, and let $\psi = \bigwedge(T' \cap \mathrm{FO}_0^{(b)})$. Then either $w'(\phi,\psi) < r+h$ or $\phi \vdash \zeta_{r+h}(\psi)$. This means that for any model $\mathbf{Y}$ of T, either $w'(\phi,\psi) < r+h$ and the root of $\mathbf{Y}$ has exactly $w'(\phi,\psi)$ sons v such that $\mathrm{Th}(\mathbf{Y}(v)) \cap \mathrm{FO}_0^{(b)} = T' \cap \mathrm{FO}_0^{(b)}$, or $w'(\phi,\psi) = r+h$ and the root of $\mathbf{Y}$ has at least $r+h$ sons v such that $\mathrm{Th}(\mathbf{Y}(v)) \cap \mathrm{FO}_0^{(b)} = T' \cap \mathrm{FO}_0^{(b)}$.

Let $\widetilde{\mu} = \mathrm{Encode}_*(\mu)$ (see Lemma 4.7) be the pushforward of μ on $\biguplus_{i=1}^h (\mathfrak{Y}_0^{(h)})^i$. For a given integer r, we define integers $a_{r,0}, a_{r,1}, \dots, a_{r,h-1}$ by

$$a_{r,h-1} = r+h, \ a_{r,h-2} = c(a_{r,h-1}, r+h), \dots, \ a_{r,0} = c(a_{r,1}, r+h).$$

Let F be the mapping defined on $\biguplus_{i=1}^h (\mathfrak{Y}_0^{(h)})^i$ by

$$F(T_0, \dots, T_i) = (T_0 \cap \mathrm{FO}_0^{(a_{r,0})}, \dots, T_i \cap \mathrm{FO}_0^{(a_{r,i})}).$$

We note that $\mathfrak{X} = F\big(\biguplus_{i=1}^h (\mathfrak{Y}_0^{(h)})^i\big)$ is a finite space, and we endow $\mathfrak{X}$ with the discrete topology. (Note that F is continuous.) We define the probability measure $\widetilde{\mu}^{(r)} = F_*(\widetilde{\mu})$ on $\mathfrak{X}$ as the pushfoward of $\widetilde{\mu}$ by F.

We will construct disjoint sets $V_{\hat{T}_0,\dots,\hat{T}_i}$ indexed by the elements $(\hat{T}_0, \dots, \hat{T}_i)$ of $\mathfrak{X}$. To construct these sets, it will be sufficient to define their cardinalities and the unary relations that apply to their elements. We proceed inductively on the length of the index tuple. As μ is pure, $\mathfrak{X}$ contains a unique 1-tuple $(\hat{T}_0)$, and we let the set $V_{\hat{T}_0}$ to be a singleton. The unique element r of $V_{\hat{T}_0}$ will be the root of the approximation tree Y_N. Hence we let $R(v)$ and for every color relation C_i we let $C_i(r)$ if $(\forall x) R(x) \to C_i(x)$ belongs to $\hat{T}_0$. Assume sets $V_{\hat{T}_0,\dots,\hat{T}_j}$ have been constructed for every $0 \leq j \leq i$ and every $(\hat{T}_0, \dots, \hat{T}_j) \in \mathfrak{X}$. Let $(\hat{T}_0, \dots, \hat{T}_{i+1}) \in \mathfrak{X}$. Then of course $(\hat{T}_0, \dots, \hat{T}_i) \in \mathfrak{X}$.

- If $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_i)\}) = 0$ and $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_{i+1})\}) = 0$ then

$$|V_{\hat{T}_0,\dots,\hat{T}_{i+1}}| = w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1}) \, |V_{\hat{T}_0,\dots,\hat{T}_i}|;$$

- If $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_i)\}) > 0$ and $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_{i+1})\}) = 0$ then (according to FMTP) $w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1}) = 0$ and we let $V_{\hat{T}_0,\dots,\hat{T}_i} = \emptyset$;
- If $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_i)\}) = 0$ and $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_{i+1})\}) > 0$ then (according to FMTP) $w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1}) = r+h$ and we let

$$|V_{\hat{T}_0,\dots,\hat{T}_{i+1}}| = \max((r+h)|V_{\hat{T}_0,\dots,\hat{T}_i}|, \lfloor \widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_{i+1})\}) N \rfloor).$$

- Otherwise $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_i)\}) > 0$ and $\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_{i+1})\}) > 0$. In this case, according to FMTP, it holds

$$w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1}) = \min\left(r+h, \frac{\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_{i+1})\})}{\widetilde{\mu}^{(r)}(\{(\hat{T}_0, \dots, \hat{T}_i)\})}\right),$$

Then, if $w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1}) < r + h$ we let

$$|V_{\hat{T}_0,\dots,\hat{T}_{i+1}}| = w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1})|V_{\hat{T}_0,\dots,\hat{T}_i}|$$

and otherwise we let

$$|V_{\hat{T}_0,\dots,\hat{T}_{i+1}}| = \max((r + h)|V_{\hat{T}_0,\dots,\hat{T}_i}|, \lfloor\widetilde{\mu}^{(r)}(\{(\hat{T}_0,\dots,\hat{T}_{i+1})\})N\rfloor).$$

The colors of the elements of $V_{\hat{T}_1,\dots,\hat{T}_i}$ are easily defined: for $v \in V_{\hat{T}_1,\dots,\hat{T}_i}$ and color relation C_i we let $C_i(v)$ if $(\forall x)R(x) \to C_i(x)$ belongs to $\hat{T}_i$.

The tree Y_N has vertex set $\bigcup V_{\hat{T}_1,\dots,\hat{T}_i}$. Each set $V_{\hat{T}_0,\dots,\hat{T}_{i+1}}$ is partitionned as equaly as possible into $|V_{\hat{T}_0,\dots,\hat{T}_i}|$ parts, each part being adjacent to a single vertex in $V_{\hat{T}_0,\dots,\hat{T}_i}$. It follows that the degree in $V_{\hat{T}_0,\dots,\hat{T}_{i+1}}$ of a vertex in $V_{\hat{T}_0,\dots,\hat{T}_i}$ lies between $\lceil|V_{\hat{T}_0,\dots,\hat{T}_{i+1}}|/|V_{\hat{T}_0,\dots,\hat{T}_i}|\rceil$ and $\lfloor|V_{\hat{T}_0,\dots,\hat{T}_{i+1}}|/|V_{\hat{T}_0,\dots,\hat{T}_i}|\rfloor$, and that (by construction and thanks to FMTP) this coincides with $w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1})$ (when $< r+h$) or is at least $w'(\bigwedge \hat{T}_i, \bigwedge \hat{T}_{i+1})$ (when $= r + h$).

For $(\hat{T}_0,\dots,\hat{T}_i) \in \mathfrak{X}$, it is easily checked that

$$\big||V_{\hat{T}_0,\dots,\hat{T}_i}| - \widetilde{\mu}(\{(\hat{T}_0,\dots,\hat{T}_i)\})N\big| \le (r + h)^i.$$

For $\phi \in \mathrm{FO}_1(\lambda)$, let $\mathcal{F}_\phi = \{F \circ \mathrm{Encode}(T) : T \in K(\phi) \cap \mathfrak{Y}_1^{(h)}\}$. Let $C = (r + h)^h|\mathfrak{X}|$. Then, by summing up the above inequality, we get

$$0 \le \left(\sum_{(\hat{T}_0,\dots,\hat{T}_i)\in\mathcal{F}_\varphi} |V_{\hat{T}_0,\dots,\hat{T}_i}|\right) - \mu(K(\varphi))N \le C.$$

In particular, if ϕ is the true statement, we get

$$N \le |Y_N| \le N + C.$$

Also, we deduce that for every $\hat{T}_0, \hat{T}_1) \in \mathfrak{X}$ and every $v_1, v_2 \in V_{\hat{T}_0,\hat{T}_1}$ it holds

$$\big||Y_N(v_1)| - |Y_N(v_2)|\big| \le C.$$

Let $\mathfrak{Z} = \{T \in \mathfrak{Y}_1^{(h)} : T \cap \mathrm{FO}_0(\lambda) = T_0\}$, let $T \in \mathfrak{Z}$, let $(T_0,\dots,T_i) = \mathrm{Encode}(T)$, and let $(\hat{T}_0,\dots,\hat{T}_i) = F(T_0,\dots,T_i)$. We now prove that if $v \in V_{\hat{T}_0,\dots,\hat{T}_i}$ and $(T_0',\dots,T_i') = \mathrm{Encode}(\mathrm{Th}(Y_N, v))$ then it holds $T_j \cap \mathrm{FO}_0^{(r+h)} = T_j' \cap \mathrm{FO}_0^{(r+h)}$ for every $1 \le j \le i$ (see Fig 4.1.4).

First note that it is sufficient to prove $T_i \cap \mathrm{FO}_0^{(r+h)} = T_i' \cap \mathrm{FO}_0^{(r+h)}$, as the other equalities follow by considering the ancestors of v. If T_i (hence T_i') is the complete theory of a single vertex tree, then by construction of $V_{\hat{T}_0,\dots,\hat{T}_i}$, it holds $T_i = T_i'$. Assume now that i is such that for every $(T_0,\dots,T_{i+1}) \in \mathrm{Encode}(T)$ with $T \in \mathfrak{Z}$, it holds $T_{i+1} \cap \mathrm{FO}_0^{(r+h)} = T_{i+1}' \cap \mathrm{FO}_0^{(r+h)}$. Let A be a model of T_i and let B be a model of T_i'. It follows from the induction that the roots of A and B have the same number of sons (up to $r + h$) with subtrees, which are $(r + h)$-equivalent to a fixed rooted tree. By an easy argument based on an Ehrefeucht-Fraïssé game, it follows that A and B are $(r + h)$-equivalent hence $T_i \cap \mathrm{FO}_0^{(r+h)} = T_i' \cap \mathrm{FO}_0^{(r+h)}$.

According to Lemma 4.4, we deduce that for $T \in \mathfrak{Z}$ and corresponding vertex $v \in V_{\hat{T}_0,\dots,\hat{T}_i}$ it holds $\mathrm{Th}(Y_N, v) \cap \mathrm{FO}_1^{(r)} = T \cap \mathrm{FO}_1^{(r)}$.

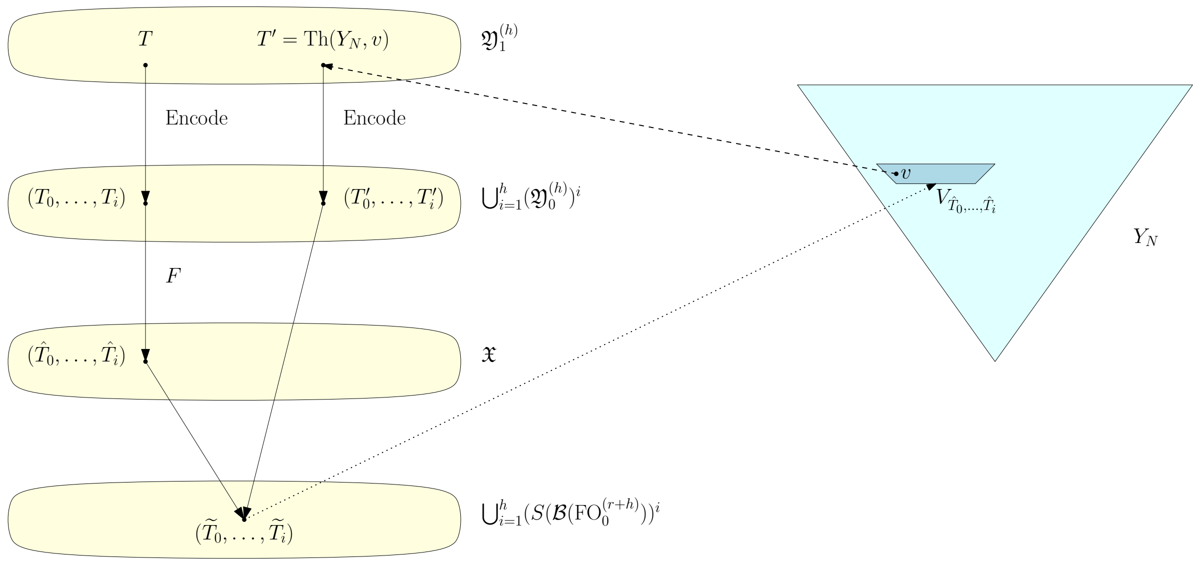

FIGURE 5. Comparing T with its approximation in Y_N

It follows that for every $\varphi \in \mathrm{FO}_1^{(r)}$ it holds

$$\langle \varphi, Y_N \rangle = \sum_{(\hat{T}_0, \ldots, \hat{T}_i) \in \mathcal{F}_\varphi} \frac{|V_{\hat{T}_0, \ldots, \hat{T}_i}|}{|Y_N|}.$$

Hence

$$\big| \langle \varphi, Y_N \rangle - \mu(K(\varphi)) \big| \leq C/N.$$

Let r_N be the root of Y_N. Define

$$\alpha_N = \max_{v \sim r_N} \frac{|Y_N(v)|}{|Y_N|}.$$

Assume $\mathbf{L}$ is a modeling with FO_1 statistics μ and root $r_{\mathbf{L}}$.

Let $(\hat{T}_0, \hat{T}_1) \in \mathfrak{X}$ (vertices in $V_{\hat{T}_0, \hat{T}_1}$ are sons of the root of Y_N). By construction, all the subtrees rooted at a vertex in $V_{\hat{T}_0, \hat{T}_1}$ have almost the same number of vertices (the difference being at most C). If $w'(\hat{T}_0, \hat{T}_1) = r + h$, it follows that for every $v \in V_{\hat{T}_0, \hat{T}_1}$ it holds $|Y_N(v)| \leq C + |Y_N|/(r+h)$, i.e. $|Y_N(v)|/|Y_N| \leq 1/(r+h) + C/N$. Otherwise, $w'(\hat{T}_0, \hat{T}_1) = k < r+h$ hence if ψ is the formula stating that the ancestor of x_1 which is a son of the root satisfies $\bigwedge \hat{T}_1$, then

$$\mu(K(\psi)) = \langle \psi, \mathbf{L} \rangle = \sum_{v \sim r_{\mathbf{L}}, \mathbf{L} \models \psi(v)} \nu_{\mathbf{L}}(\mathbf{L}(v)) \leq k \sup_{v \sim r_{\mathbf{L}}} \nu_{\mathbf{L}}(\mathbf{L}(v)).$$

Also

$$\begin{aligned} \mu(K(\psi)) + \frac{C}{N} \geq \langle \psi, Y_N \rangle &= \sum_{(\hat{T}_0, \ldots, \hat{T}_i) \in \mathcal{F}_\psi} \frac{|V_{\hat{T}_0, \ldots, \hat{T}_i}|}{|Y_N|} \\ &= \sum_{v \sim r_N, Y_n \models \psi(v)} \frac{|Y_N(v)|}{|Y_N|} \end{aligned}$$

hence $\mu(K(\psi)) + \frac{C}{N} \geq \max_{v\sim r_N, Y_n\models\psi(v)} \frac{|Y_N(v)|}{|Y_N|}$ if $k=1$, and otherwise

$$\begin{aligned}\mu(K(\psi)) + \frac{C}{N} &\geq k \max_{v\sim r_N, Y_n\models\psi(v)} \frac{|Y_N(v)| - C}{|Y_N|}\\ &\geq k \max_{v\sim r_N, Y_n\models\psi(v)} \frac{|Y_N(v)|}{|Y_N|} - C/N.\end{aligned}$$

Hence, considering the case $k=1$ and the case $k \geq 2$ (where $2C/k \leq C$) we get

$$\max_{v\sim r_N, Y_n\models\psi(v)} \frac{|Y_N(v)|}{|Y_N|} \leq \sup_{v\sim r_{\mathbf{L}}} \nu_{\mathbf{L}}(\mathbf{L}(v)) + C/N.$$

And we deduce

$$\alpha_N \leq \max\Bigl(\frac{1}{r+h}, \sup_{v\sim r_{\mathbf{L}}} \nu_{\mathbf{L}}(\mathbf{L}(v))\Bigr) + C/N.$$

□

We get the following two inverse results:

THEOREM 4.23. *A measure μ on $S(\mathcal{B}(\mathrm{FO}_1))$ is the weak limit of a sequence of measures μ_{Y_n} associated to an FO_1-convergent sequence $(Y_n)_{n\in\mathbb{N}}$ of finite colored rooted trees with height at most h (i.e. of finite $Y_n \in \mathcal{Y}^{(h)}$) if and only if*

- *μ is pure and its complete theory belongs to $\mathfrak{Y}_0^{(h)}$,*
- *μ satisfies the FMTP.*

PROOF. Assume that $(Y_n)_{n\in\mathbb{N}}$ is an FO_1-convergent sequence of finite $Y_n \in \mathcal{Y}^{(h)}$, and that $\mu_{Y_N} \Rightarrow \mu$. According to Remark 2.10, μ is pure. As $(Y_n)_{n\in\mathbb{N}}$ is elementarily convergent, the complete theory of μ is the complete theory of the elementary limit of $(Y_n)_{n\in\mathbb{N}}$. Also, μ satisfies the FMTP (see Section 2.1.4).

Conversely, assume μ is pure, that its complete theory belongs to $\mathfrak{Y}_0^{(h)}$, and that it satisfies the FMTP. According to Lemma 4.22 we can construct a sequence $(Y_n)_{n\in\mathbb{N}}$ of finite $Y_n \in \mathcal{Y}^{(h)}$ (considering for instance $r=n$ and $N = 10^C$ where C is the constant defined from r, h, c) such that for every formula $\phi \in \mathrm{FO}_1(\lambda)$ it holds $\bigl|\langle\phi, Y_n\rangle - \mu(K(\phi))\bigr| \to 0$ as $n\to\infty$, i.e. $\mu_{Y_n} \Rightarrow \mu$. □

and we deduce

THEOREM 4.24. *A modeling $\mathbf{L}$ is the FO_1-limit of an FO_1-convergent sequence $(Y_n)_{n\in\mathbb{N}}$ of finite colored rooted trees with height at most h (i.e. of finite $Y_n \in \mathcal{Y}^{(h)}$) if and only if*

- *$\mathbf{L}$ is a colored rooted tree with height at most h (i.e. $\mathbf{L} \in \mathcal{Y}^{(h)}$),*
- *$\mathbf{L}$ satisfies the FMTP.*

PROOF. That an FO_1-convergent sequence of finite rooted colored trees $Y_n \in \mathcal{Y}^{(h)}$ has a modeling FO_1-limit is a direct consequence of Theorem 4.18. That $\mathbf{L}$ satisfies the FMTP is immediate (as the associate measure $\mu = {}^1_{\mathbf{L}}\mathsf{Tp}_*(\nu_{\mathbf{L}})$ does).

Conversely, that a colored rooted tree modeling $\mathbf{L} \in \mathcal{Y}^{(h)}$ that satisfies the FMTP is the FO_1-limit of a sequence of finite rooted colored trees $Y_n \in \mathcal{Y}^{(h)}$ is a direct consequence of Theorem 4.23. □

4.2. FO-limits of Colored Rooted Trees with Bounded Height

In this section we explicitly define modeling FO-limits for FO-convergent sequences of colored rooted trees with bounded height.

We first sketch our method.

We consider the signature λ^+, which is the signature λ augmented by a new unary relation P. Particular λ^+-structures are colored rooted forests with a *principal* connected component, whose root will be marked by relation P instead of R (no other vertex gets P). The class of colored rooted forests with a principal connected component and height at most h will be denoted by $\mathcal{F}^{(h)}$.

We consider three basic interpretation schemes:

(1) $\mathsf{I}_{Y\to F}$ is a basic interpretation scheme of λ^+-structures in λ-structures defined as follows: for every λ-structure $\mathbf{A}$, the domain of $\mathsf{I}_{Y\to F}(\mathbf{A})$ is the same as the domain of $\mathbf{A}$, and it holds (for every $x, y \in A$):

$$\begin{aligned}
\mathsf{I}_{Y\to F}(\mathbf{A}) \models x \sim y \quad &\iff \quad \mathbf{A} \models (x \sim y) \wedge \neg R(x) \wedge \neg R(y)\\
\mathsf{I}_{Y\to F}(\mathbf{A}) \models R(x) \quad &\iff \quad \mathbf{A} \models (\exists z)\ R(z) \wedge (z \sim x)\\
\mathsf{I}_{Y\to F}(\mathbf{A}) \models P(x) \quad &\iff \quad \mathbf{A} \models R(x)
\end{aligned}$$

In particular, $\mathsf{I}_{Y\to F}$ maps a colored rooted tree $Y \in \mathcal{Y}^{(h)}$ into a colored rooted forest $\mathsf{I}_{Y\to F}(\mathbf{A})(Y) \in \mathcal{F}^{(h-1)}$, formed by the subtrees rooted at the sons of the former root and a single vertex rooted principal component (the former root);

(2) $\mathsf{I}_{F\to Y}$ is a basic interpretation scheme of λ-structures in λ^+-structures defined as follows: for every λ^+-structure $\mathbf{A}$, the domain of $\mathsf{I}_{F\to Y}(\mathbf{A})$ is the same as the domain of $\mathbf{A}$, and it holds (for every $x, y \in A$):

$$\begin{aligned}
\mathsf{I}_{F\to A}(\mathbf{A}) \models x \sim y \quad &\iff \quad \mathbf{A} \models (x \sim y) \vee R(x) \wedge P(y) \vee R(y) \wedge P(x)\\
\mathsf{I}_{F\to A}(\mathbf{A}) \models R(x) \quad &\iff \quad \mathbf{A} \models P(x)
\end{aligned}$$

In particular, $\mathsf{I}_{F\to Y}$ maps a colored rooted forest $F \in \mathcal{F}^{(h)}$ into a colored rooted tree $\mathsf{I}_{F\to Y}(F) \in \mathcal{Y}^{(h+1)}$ by making each non principal root a son of the principal root;

(3) $\mathsf{I}_{R\to P}$ is a basic interpretation scheme of λ^+-structures in λ-structures defined as follows: for every λ^+-structure $\mathbf{A}$, the domain of $\mathsf{I}_{R\to P}(\mathbf{A})$ is the same as the domain of $\mathbf{A}$, adjacencies are the same in $\mathbf{A}$ and $\mathsf{I}_{R\to P}(\mathbf{A})$, no element of $\mathsf{I}_{R\to P}(\mathbf{A})$ is in R, and for every $x \in A$ it holds

$$I_{R\to P}(\mathbf{A}) \models P(x) \quad \iff \quad \mathbf{A} \models R(x).$$

(Roughly speaking, the relation R becomes the relation P.) In particular, $\mathsf{I}_{R\to P}$ maps a colored rooted tree $Y \in \mathcal{Y}^{(h)}$ into a colored rooted forest $\mathsf{I}_{R\to P}(Y) \in \mathcal{F}^{(h)}$ having a single (principal) component.

We now outline our proof strategy. Let $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ be an FO-convergent sequence of finite rooted colored trees ($\mathbf{Y}_n \in \mathcal{Y}^{(h)}$) such that $\lim_{n\to\infty} |Y_n| = \infty$.

For each n, $\mathsf{I}_{Y\to F}(\mathbf{Y}_n)$ is a forest $\mathbf{F}_n$, and $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is an FO-convergent sequence. According to the Comb Structure Theorem, there exists a countable set $(\mathbf{Y}_{n,i})_{n\in\mathbb{N}}$ of FO-convergent sequences of colored rooted trees $\mathbf{Y}_{n,i} \in \mathcal{Y}^{(h)}$ and a FO-convergent sequence $(\mathbf{R}_n)_{n\in\mathbb{N}}$ of residues $\mathbf{R}_n \in \mathcal{F}^{(h)}$, which are special colored rooted forests (as the isolated principal root obviously belongs to $\mathbf{R}_n$), so that

- the sequences $(\mathbf{Y}_{n,i})_{n\in\mathbb{N}}$ and the sequence $(\mathbf{R}_n)_{n\in\mathbb{N}}$ form a uniformly convergent family of sequences;
- for each $n \in \mathbb{N}$ it holds $\mathsf{I}_{Y\to F}(\mathbf{Y}_n) = \mathbf{R}_n \cup \bigcup_{i\in I} \mathbf{Y}_{n,i}$.

If the limit spectrum of $(\mathsf{I}_{Y\to F}(\mathbf{Y}_n))_{n\in\mathbb{N}}$ is empty (i.e. $I = \emptyset$), the sequence $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ of colored rooted trees is called *residual*, and in this case we deduce directly that a residual sequence of colored rooted trees admit a modeling FO-limit from our results on FO_1-convergent sequences.

Otherwise, we proceed by induction over the height bound h. Denote by $(\mathsf{sp}_i)_{i\in I}$ the limit spectrum of $(\mathsf{I}_{Y\to F}(\mathbf{Y}_n))_{n\in\mathbb{N}}$, let $\mathsf{sp}_0 = 1 - \sum_{i\in I} \mathsf{sp}_i$, and let $\mathbf{Y}_{n,0} = \mathsf{I}_{R\to P} \circ \mathsf{I}_{F\to Y}(\mathbf{R}_n)$. As $(\mathsf{I}_{F\to Y}(\mathbf{R}_n))_{n\in\mathbb{N}}$ is residual, $(\mathbf{Y}_{n,0})_{n\in\mathbb{N}}$ has a modeling FO-limit $\widetilde{\mathbf{Y}}_0$. By induction, each $(\mathbf{Y}_{n,i})_{n\in\mathbb{N}}$ has a modeling FO-limit $\widetilde{\mathbf{Y}}_i$. As $\mathbf{Y}_n = \mathsf{I}_{F\to Y}(\bigcup_{i\in I\cup\{0\}} \mathbf{Y}_{n,i})$, we deduce (using uniform elementary convergence) that $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ has modeling FO-limit $\mathsf{I}_{F\to Y}(\coprod_{i\in I\cup\{0\}}(\widetilde{\mathbf{Y}}_i, \mathsf{sp}_i))$.

This finishes the outline of our construction. Now we provide details.

4.2.1. The Modeling FO-limit of Residual Sequences. We start by a formal definition of residual sequences of colored rooted trees.

DEFINITION 4.25. Let $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ be a sequence of finite colored rooted trees, let N_n be the set of all sons of the root of $\mathbf{Y}_n$, and let $\mathbf{Y}_n(v)$ denote (for $v \in Y_n$) the subtree of $\mathbf{Y}_n$ rooted at v.

The sequence $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is *residual* if

$$\limsup_{n\to\infty} \max_{v\in N_n} \frac{|\mathbf{Y}_n(v)|}{|\mathbf{Y}_n|} = 0.$$

We extend this definition to single infinite modelings.

DEFINITION 4.26. A modeling colored rooted tree $\widetilde{\mathbf{Y}}$ with height at most h is *residual* if, denoting by N the neighbour set of the root, it holds

$$\sup_{v\in N} \nu_{\widetilde{\mathbf{Y}}}(\widetilde{\mathbf{Y}}(v)) = 0.$$

Note that the above definition makes sense as belonging to a same $\widetilde{\mathbf{Y}}(v)$ (for some $v \in N$) is first-order definable hence, as $\widetilde{\mathbf{Y}}$ is a relational sample space, each $\widetilde{\mathbf{Y}}(v)$ is $\Sigma_{\widetilde{\mathbf{Y}}}$-measurable.

We first prove that for a modeling colored rooted tree to be a modeling FO-limit of a residual sequence $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ of rooted colored trees with bounded height, it is sufficient that it is a modeling FO_1-limit of the sequence.

LEMMA 4.27. *Assume $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is a residual FO_1-convergent sequence of finite rooted colored trees with bounded height with residual modeling FO_1-limit $\widetilde{\mathbf{Y}}$.*

Then $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is FO-convergent and has modeling FO-limit $\widetilde{\mathbf{Y}}$.

PROOF. Let h be a bound on the height of the rooted trees $\mathbf{Y}_n$. Let $\mathbf{F}_n = I_{Y\to F}(\mathbf{Y}_n)$. Let ϖ be the formula asserting $\mathrm{dist}(x_1, x_2) \le 2h$. Then $\mathbf{F}_n \models \varpi(u, v)$ holds if and only if u and v belong to a same connected component of $\mathbf{F}_n$. According to Lemma 3.47, we get that $(\mathbf{F}_n)_{n\in\mathbb{N}}$ is $\mathrm{FO}^{\mathrm{local}}$ convergent. As it is also FO_0-convergent, it is FO convergent (according to Theorem 2.21). As $\mathbf{Y}_n = I_{F\to Y}(\mathbf{F}_n)$, we deduce that $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is FO-convergent.

That $\widetilde{\mathbf{Y}}$ is a modeling FO-limit of $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ then follows from Theorem 3.17. $\square$

LEMMA 4.28. *Let* $\mathbf{Y}_n$ *be a residual* FO_1*-convergent sequence of colored rooted trees with height at most* h*, let* μ *be the limit measure of* $\mu_{\mathbf{Y}_n}$ *on* $\mathfrak{T}_1^{(h)}$*, and let* $\widetilde{\mathbf{Y}}$ *be the connected component of* $\mathbb{Y}_h$ *containing the support of* ν*. Then* $\widetilde{\mathbf{Y}}$*, equipped with the probability measure* $\nu_{\widetilde{\mathbf{Y}}} = \nu$*, is a modeling* FO*-limit of* $(\mathbf{Y}_n)_{n\in\mathbb{N}}$*.*

PROOF. That $\widetilde{\mathbf{Y}}$ is a residual FO_1-modeling limit of $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ is a consequence of Theorem 4.18. That it is then an FO-modeling limit of $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ follows from Lemma 4.27 □

4.2.2. The Modeling FO-Limit of a Sequence of Rooted Trees. For an intuition of how the structure of a modeling FO-limit of a sequence of colored rooted trees with height at most h could look like, consider a modeling rooted colored tree $\mathbf{Y}$. Obviously, the $\mathbf{Y}$ contains two kind of vertices: the *heavy* vertices v such that the subtree $\mathbf{Y}(v)$ of $\mathbf{Y}$ rooted at v has positive $\nu_{\mathbf{Y}}$-measure and the *light* vertices for which $\mathbf{Y}(v)$ has zero $\nu_{\mathbf{Y}}$-measure. It is then immediate that heavy vertices of $\mathbf{Y}$ induce a countable rooted subtree with same root as $\mathbf{Y}$.

This suggest the following definitions.

DEFINITION 4.29. A *rooted skeleton* is a countable rooted tree $\mathbf{S}$ together with a *mass function* $m : S \to (0, 1]$ such that $m(r) = 1$ (r is the root of $\mathbf{S}$) and for every non-leaf vertex $v \in S$ it holds

$$m(v) \geq \sum_{u \text{ son of } v} m(u).$$

DEFINITION 4.30. Let $(\mathbf{S}, m)$ be a rooted skeleton, let S_0 be the subset of S with vertices v such that $m(v) = \sum_{u \text{ son of } v} m(u)$, let $(\mathbf{R}_v)_{v\in S\setminus S_0}$ be a countable sequence of non-empty residual λ-modeling indexed by $S \setminus S_0$, and let $(\mathbf{R}_v)_{v\in S_0}$ be a countable sequence of non empty countable colored rooted trees indexed by S_0. The *grafting* of $(\mathbf{R}_v)_{v\in S\setminus S_0}$ and $(\mathbf{R}_v)_{v\in S_0}$ on $(\mathbf{S}, m)$ is the modeling $\mathbf{Y}$ defined as follows: As a graph, $\mathbf{Y}$ is obtained by taking the disjoint union of $\mathbf{S}$ with the colored rooted trees $\mathbf{R}_v$ and then identifying $v \in S$ with the root of $\mathbf{R}_v$ (see Fig. 6). The sigma algebra $\Sigma_{\mathbf{Y}}$ is defined as

$$\Sigma_{\mathbf{Y}} = \Big\{ \bigcup_{v\in S\setminus S_0} M_v \cup \bigcup_{v\in S_0} M'_v :\ M_v \in \Sigma_{\mathbf{R}_v}, M'_v \subseteq R_v \Big\}$$

and the measure $\nu_{\mathbf{Y}}(M)$ of $M \in \Sigma$ is defined by

$$\nu_{\mathbf{Y}}(M) = \sum_{v\in S\setminus S_0} \Big(m(v) - \sum_{u \text{ son of } v} m(u) \Big) \nu_{\mathbf{R}_v}(M_v),$$

where $M = \bigcup_{v\in S\setminus S_0} M_v \cup \bigcup_{v\in S} M'_v$ with $M_v \in \Sigma_{\mathbf{R}_v}$ and $M'_v \subseteq R_v$.

LEMMA 4.31. *Let* $\mathbf{Y}$ *be obtained by grafting countable sequence of non-empty modeling colored rooted trees* $\mathbf{R}_v$ *on a rooted skeleton* $(\mathbf{S}, m)$*. Then* $\mathbf{Y}$ *is a modeling.*

PROOF. We prove the statement by induction over the height of the rooted skeleton. The statement obviously holds if $\mathbf{S}$ is a single vertex rooted tree (that is if $\mathrm{height}(\mathbf{S}) = 1$). Assume that the statement holds for rooted skeletons with height at most h, and let $(\mathbf{S}, m)$ be a rooted skeleton with height $h + 1$.

Let s_0 be the root of $\mathbf{S}$ and let $\{s_i : i \in I \subseteq \mathbb{N}\}$ be the set of the sons of s_0 in $\mathbf{S}$. For $i \in I$, $\mathbf{Y}_i = \mathbf{Y}(s_i)$ be the subtree of $\mathbf{Y}$ rooted at s_i, let $\mathsf{sp}_i = \sum_{x\in Y_i} m(x)$,

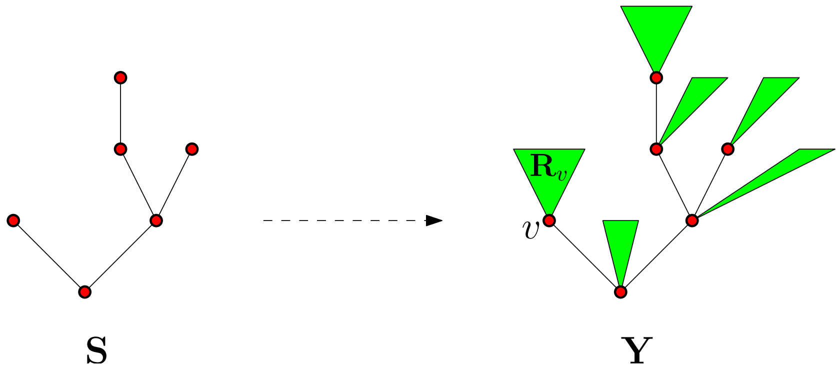

FIGURE 6. Grafting of trees on a skeleton

and let m_i be the mass function on $\mathbf{S}_i$ defined by $m_i(v) = m(v)/\mathsf{sp}_i$. Also, let $\mathsf{sp}_0 = 1 - \sum_{i\in I} \mathsf{sp}_i$.

For each $i \in I \cup \{0\}$, if $\mathsf{sp}_i = 0$ (in which case $\mathbf{R}_{s_i}$ is only assumed to be a relational sample space) we turn $\mathbf{R}_{s_i}$ into a modeling by defining a probability measure on $\mathbf{R}_{s_i}$ concentrated on s_i.

For $i \in I$, let $\mathbf{Y}_i$ be obtained by grafting the $\mathbf{R}_v$ on $(\mathbf{S}_i, m_i)$ (for $v \in S_i$), and let $\mathbf{Y}_0$ be the λ^+-modeling consisting in a rooted colored forest with single (principal) component $\mathbf{R}_{s_0}$ (that is: $\mathbf{Y}_0 = \mathsf{I}_{R\to P}(\mathbf{R}_{s_0})$). According to Lemma 3.11, $\mathbf{Y}_0$ is a modeling, and by induction hypothesis each $\mathbf{Y}_i$ $(i \in I)$ is a modeling. According to Lemma 3.21, it follows that $\mathbf{F} = \amalg_{i\in I\cup\{0\}}(\mathbf{Y}_i, \mathsf{sp}_i)$ is a modeling. Hence, according to Lemma 3.11, $\mathbf{Y} = \mathsf{I}_{F\to Y}(\mathbf{F})$ is a modeling. $\square$

Our main theorem is the following.

THEOREM 4.32. *Let $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ be an* FO*-convergent sequence of finite colored rooted trees with height at most h.*

Then there exists a skeleton $(\mathbf{S}, m)$ and a family $(\mathbf{R}_v)_{v\in S}$ — where $\mathbf{R}_v$ is (isomorphic to) a connected component of $\mathbb{Y}_h$, $\Sigma_{\mathbf{R}_v}$ is the induced σ-algebra on R_v — with the property that the grafting $\mathbf{Y}$ of the $\mathbf{R}_v$ on $(\mathbf{S}, m)$ is a modeling FO*-limit of the sequence $(\mathbf{Y}_n)_{n\in\mathbb{N}}$.*

PROOF. First notice that the statement obviously holds if $\lim_{n\to\infty} |Y_n| < \infty$ as then the sequence is eventually constant to a finite colored rooted tree $\mathbf{Y}$: we can let $\mathbf{S}$ be $\mathbf{Y}$ (without the colors), m be the uniform weight ($m(v) = 1/|Y|$), and $\mathbf{R}_v$ be single vertex rooted tree whose root's color is the color of v in $\mathbf{Y}$. So, we can assume that $\lim_{n\to\infty} |Y_n| = \infty$.

We prove the statement by induction over the height bound h. For $h = 1$, each $\mathbf{Y}_n$ is a single vertex colored rooted tree, and the statement obviously holds.

Assume that the statements holds for $h = h_0 - 1 \geq 1$ and let finite colored rooted trees with height at most h_0. Let $\mathbf{F}_n = \mathsf{I}_{Y\to F}(\mathbf{Y}_n)$. Then $(\mathbf{F}_n)_{n\in\mathbb{N}}$ is FO-convergent (according to Lemma 3.11). According to the Comb Structure Theorem, there exists countably many convergent sequences $(\mathbf{Y}_{n,i})_{n\in\mathbb{N}}$ of colored rooted trees (for $i \in I$) and an FO-convergent sequence $(\mathbf{R}_n)_{n\in\mathbb{N}}$ of special rooted forests forming a uniformly convergent family of sequences, such that $\mathsf{I}_{Y\to F}(\mathbf{Y}_n) = \mathbf{R}_n \cup \bigcup_{i\in I} \mathbf{Y}_{n,i}$.

If the limit spectrum of $(\mathsf{I}_{Y\to F}(\mathbf{Y}_n))_{n\in\mathbb{N}}$ is empty (i.e. $I = \emptyset$), the sequence $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ of colored rooted trees is *residual*, and the result follows from Lemma 4.28.

Otherwise, let $(\mathsf{sp}_i)_{i\in I}$ the limit spectrum of $(\mathsf{I}_{Y\to F}(\mathbf{Y}_n))_{n\in\mathbb{N}}$, let $\mathsf{sp}_0 = 1-\sum_{i\in I}\mathsf{sp}_i$, and let $\mathbf{Y}_{n,0} = \mathsf{I}_{R\to P}\circ\mathsf{I}_{F\to Y}(\mathbf{R}_n)$. If $\mathsf{sp}=0$ then there is a connected component $\widetilde{\mathbf{Y}}_0$ of $\mathbb{Y}_h$ that is an elementary limit of $(\mathbf{Y}_{n,0})_{n\in\mathbb{N}}$; Otherwise, as $(\mathsf{I}_{F\to Y}(\mathbf{R}_n))_{n\in\mathbb{N}}$ is residual, $(\mathbf{Y}_{n,0})_{n\in\mathbb{N}}$ has, according to Lemma 4.28, a modeling FO-limit $\widetilde{\mathbf{Y}}_0$. By induction, each $(\mathbf{Y}_{n,i})_{n\in\mathbb{N}}$ has a modeling FO-limit $\widetilde{\mathbf{Y}}_i$. As $\mathbf{Y}_n = \mathsf{I}_{F\to Y}(\bigcup_{i\in I\cup\{0\}}\mathbf{Y}_{n,i})$, we deduce, by Corollary 3.4, Lemma 3.27, Theorem 2.21, and Lemma 3.11, that $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ has modeling FO-limit $\mathsf{I}_{F\to Y}(\coprod_{i\in I\cup\{0\}}(\widetilde{\mathbf{Y}}_i,\mathsf{sp}_i))$. □

So, in the case of colored rooted trees with bounded height, we have constructed an explicit relational sample space that allows to pullback the limit measure μ defined on the Stone space $S(\mathcal{B}(\mathrm{FO}))$.

4.2.3. Inverse theorem for FO-limits of Colored Rooted Trees with Bounded Height. Recall that for λ-modelings $\mathbf{A}$ and $\mathbf{B}$, and $p,r\in\mathbb{N}$ we defined

$$\|\mathbf{A}-\mathbf{B}\|_{p,r}^{\mathrm{local}} = \sup\{|\langle\phi,\mathbf{A}\rangle-\langle\phi,\mathbf{B}\rangle|:\ \phi\in\mathrm{FO}_p^{\mathrm{local}}(\lambda),\mathrm{qrank}(\phi)\le r\}.$$

LEMMA 4.33. *Let $\mathbf{L}\in\mathcal{Y}^{(h)}$ (with root $r_{\mathbf{L}}$) be a colored rooted tree modeling that satisfies the FMTP, let $p,r\in\mathbb{N}$, and let $\epsilon>0$. Then there exist $C_0 = C_0(\lambda,p,r,\epsilon), N_0 = N_0(\lambda,p,r,\epsilon)$ such that for every $N\ge N_0$ there exists there exists a finite colored rooted tree $Y\in\mathcal{Y}^{(h)}$ such that it holds $N\le|Y|\le N+C_0$, $Y\equiv^r\mathbf{L}$, and*

$$\|Y-\mathbf{L}\|_{p,r}^{\mathrm{local}} < \max(\epsilon, 2\sup_{v\sim r_{\mathbf{L}}}\nu_{\mathbf{L}}(\mathbf{L}(v))).$$

PROOF. Without loss of generality we can assume $\epsilon\ge 2\sup_{v\sim r_{\mathbf{L}}}\nu_{\mathbf{L}}(\mathbf{L}(v))$. Let $r'=\max(r,4c_{r,p}/\epsilon)$, where $c_{r,p}$ is the constant introduced in Lemma 3.23..

According to Lemma 4.22, there is $C_0 = C(\lambda,r')$ (hence C_0 depends on λ,p,r, and ϵ) such that for every $N\in\mathbb{N}$ there exists $Y\in\mathcal{Y}^{(h)}$ with the following properties:

(1) $N\le|Y|\le N+C_0$;

(2) for every $\varphi\in\mathrm{FO}_1$ with quantifier rank at most r' it holds

$$\big|\langle\varphi,Y\rangle-\langle\varphi,\mathbf{L}\rangle\big|\le C_0/N.$$

(In particular $Y\equiv^{r'}\mathbf{L}$ as $N>C_0$.)

(3) we have

$$\max_{v\sim r_N}\frac{|Y(v)|}{|Y|}\le\max\Big(\frac{1}{r'+h},\sup_{v\sim r_{\mathbf{L}}}\nu_{\mathbf{L}}(\mathbf{L}(v))\Big)+C_0/N.$$

Let $N_0 = 4c_{r,p}C(\lambda,r')/\epsilon$ and assume $N\ge N_0$.

Let $F=\mathsf{I}_{Y\to F}(Y)$ and $\mathbf{A}=\mathsf{I}_{Y\to F}(\mathbf{L})$. Let $F_i, i\in\Gamma_F$ and $\mathbf{A}_i, i\in\Gamma_{\mathbf{A}}$ be the connected components of F and $\mathbf{A}$. Then

$$\max_{i\in\Gamma_F}\frac{|F_i|}{|F|}\le\max\Big(\frac{1}{r'+h},\sup_{i\in\Gamma_{\mathbf{L}}}\nu_{\mathbf{A}}(\mathbf{A}_i)\Big)+C_0/N<\frac{\epsilon}{2c_{r,p}}.$$

As $\mathsf{I}_{Y\to F}$ is a quantifier-free interpretation, for every $\varphi\in\mathrm{FO}_1$ with quantifier rank at most r it holds

$$\big|\langle\varphi,F\rangle-\langle\varphi,\mathbf{A}\rangle\big|\le C_0/N\le\frac{\epsilon}{4c_{r,p}}.$$

In particular we have $\|F-\mathbf{A}\|_{1,r}^{\mathrm{local}}\le\frac{\epsilon}{4c_{r,p}}$. According to Lemma 3.24, it holds

$$\|F-\mathbf{A}\|_{p,r}^{\mathrm{local}}<c_{r,p}\left(\max_{i\in\Gamma_F}\frac{|F_i|}{|F|}+\sup_{i\in\Gamma_{\mathbf{L}}}\nu_{\mathbf{A}}(\mathbf{A}_i)+\|F-\mathbf{A}\|_{1,r}^{\mathrm{local}}\right)<\epsilon.$$

and it follows that $\|Y - \mathbf{L}\|_{p,r}^{\text{local}} < \epsilon$. $\square$

LEMMA 4.34. *Let $\mathbf{L} \in \mathcal{Y}^{(h)}$ be an infinite colored rooted tree modeling that satisfies the FMTP, let $p, r \in \mathbb{N}$ and let $\epsilon > 0$. Then there exist constants C_h, N_h (depending on λ, p, r, ϵ) such that for every $N \geq N_h$ there is a finite colored rooted tree $Y_\epsilon \in \mathcal{Y}^{(h)}$ such that $N \leq |Y_\epsilon| \leq N + C_h$, $Y_\epsilon \equiv^r \mathbf{L}$, and $\|\mathbf{L} - Y_\epsilon\|_{p,r}^{\text{local}} < \epsilon$.*

PROOF. Let $\alpha = \epsilon^2/(2(3c_{r,p})^h)$, where $c_{r,p}$ is the constant which appears in Lemma 3.23. A vertex $v \in L$ is α-heavy if either v is the root $r_{\mathbf{L}}$ of $\mathbf{L}$, or the father u of v in $\mathbf{L}$ is α-heavy and $\nu_{\mathbf{L}}(\mathbf{L}(v)) > \alpha\nu_{\mathbf{L}}(\mathbf{L}(u))$. The α-heavy vertices of $\mathbf{L}$ form a finite subtree S rooted at $r_{\mathbf{L}}$ (each node v of S has at most $1/\alpha$ sons).

We prove by induction on the height t of S that — assuming $\alpha \leq \epsilon/2$ — there exist constants C_{t-1}, N_{t-1} (depending on λ, p, r, ϵ) such that for every $N \geq N_{t-1}$ there is a finite colored rooted tree $Y_\epsilon \in \mathcal{Y}^{(h)}$ such that $N \leq |Y_\epsilon| \leq N + C_{t-1}$, $Y_\epsilon \equiv^r \mathbf{L}$, and $\|\mathbf{L} - Y_\epsilon\|_{p,r}^{\text{local}} < \epsilon$.

If $t = 1$ (i.e. $r_{\mathbf{L}}$ is the only α-heavy vertex) then

$$\sup_{v \sim r_{\mathbf{L}}} \nu_{\mathbf{L}}(\mathbf{L}(v)) < \alpha.$$

Hence, according to Lemma 4.33, there exists N_0, C_0 (depending on λ, r, p, and ϵ) such that for every $N \geq N_0$ there is a finite colored rooted tree $Y \in \mathcal{Y}^{(h)}$ such that $N \leq |Y| \leq N + C_0$, $Y \equiv^r \mathbf{L}$, and

$$\|Y - \mathbf{L}\|_{p,r}^{\text{local}} < \max(\epsilon, 2 \sup_{v \sim r_{\mathbf{L}}} \nu_{\mathbf{L}}(\mathbf{L}(v))) = \epsilon.$$

Now assume that the statement we want to prove by induction holds when S has height at most $t \geq 1$, and let $\mathbf{L}$ be such that the associated subtree S of α-heavy vertices has height $t+1$. Let $v_1, \dots, v_k$ (where k is at most $1/\alpha$) be the α-heavy sons of the root $r_{\mathbf{L}}$ of $\mathbf{L}$, let $\mathbf{L}_i$ be the relational sample space defined by $\mathbf{L}_i = \mathbf{L}(v_i)$ for $1 \leq i \leq k$, and let $\mathbf{L}_0$ be the colored rooted tree relational sample space obtained by removing all the subtrees $\mathbf{L}_i$ from $\mathbf{L}$. Each L_i is measurable in $\mathbf{L}$. Let $a_i = \nu_{\mathbf{L}}(\mathbf{L}_i)$, and let

$$\begin{aligned}
\epsilon' &= \frac{\epsilon}{3c_{r,p}}\\
C_t(\lambda, p, r, \epsilon) &= \max\left(\frac{C_{t-1}(\lambda, p, r, \epsilon')}{\alpha}, C_0(\lambda, p, r, \epsilon'/3c_{r,p})\right)\\
N_t(\lambda, p, r, \epsilon) &= \max\left(\frac{N_{t-1}(\lambda, p, r, \epsilon')}{\epsilon'}, \frac{C_{t-1}(\lambda, p, r, \epsilon')}{\alpha\epsilon'}, N_0(\lambda, p, r, \epsilon'/3c_{r,p})\right).
\end{aligned}$$

(Nothe that we do not change α.)

Assume $a_0 \geq \epsilon'$. Let $\hat{\mathbf{L}}_{\mathbf{i}}$ be the modeling with relational sample space $\mathbf{L}_i$ and probability measure $\nu_{\hat{\mathbf{L}}_{\mathbf{i}}}$ which is $a_i^{-1}\nu_{\mathbf{L}}|L_i$, where $\nu_{\mathbf{L}}|L_i$ stands for the restriction of $\nu_{\mathbf{L}}$ to L_i. Let S_i be the rooted subtree of α-heavy vertices of $\hat{\mathbf{L}}_{\mathbf{i}}$. Clearly, if $1 \leq i \leq k$, then $S_i = S(v_i)$ (as we did not change α) thus S_i has height at most t. Let $\mathbf{F} \in \mathcal{F}^{(h)}$ be the forest defined from $\mathbf{F} = \coprod_{i=0}^k (\hat{\mathbf{L}}_{\mathbf{i}}, a_i)$ by making the component $\hat{\mathbf{L}}_{\mathbf{0}}$ special. It is clear that $\mathbf{L} = \mathsf{l}_{F \to Y}(\mathbf{F})$. For every $N \geq N_t(\lambda, p, r, \epsilon) \geq N_{t-1}(\lambda, p, r, \epsilon')/\epsilon'$ there exist, by induction, $Y_1, \dots, Y_k$ such that $a_i N \leq Y_i \leq a_i N + C_{t-1}(\lambda, p, r, \epsilon')$, $Y_i \equiv^r \hat{\mathbf{L}}_{\mathbf{i}}$, and $\|Y_i - \hat{\mathbf{L}}_{\mathbf{i}}\|_{p,r}^{\text{local}} < \epsilon'$. As the induction step is carried on at most h

times, it will always hold that $\alpha \leq \epsilon'^2/2$ hence

$$\sup_{v\sim r_{\hat{\mathbf{L}}_\mathbf{0}}} \nu_{\hat{\mathbf{L}}_\mathbf{0}}(\hat{\mathbf{L}}_\mathbf{0}(v))) \leq \frac{1}{\epsilon'} \sup_{v\sim r_{\mathbf{L}_\mathbf{0}}} \nu_{\mathbf{L}}(\mathbf{L}_\mathbf{0}(v))) \leq \frac{\alpha}{\epsilon'} \leq \epsilon'/2.$$

Also, according to Lemma 4.33, for every $N \geq N_{t-1}(\lambda,p,r,\epsilon') \geq N_0(\lambda,p,r,\epsilon')$ there is a finite colored rooted tree $Y_0 \in \mathcal{Y}^{(h)}$ such that $N \leq |Y_0| \leq N + C_0(\lambda,p,r,\epsilon') \leq N + C_{t-1}(\lambda,p,r,\epsilon')$, $Y' \equiv^r \hat{\mathbf{L}}_\mathbf{0}$, and

$$\|Y_0 - \hat{\mathbf{L}}_\mathbf{0}\|_{p,r}^{\mathrm{local}} < \max(\epsilon', 2\sup_{v\sim r_{\hat{\mathbf{L}}_\mathbf{0}}} \nu_{\hat{\mathbf{L}}_\mathbf{0}}(\hat{\mathbf{L}}_\mathbf{0}(v))) = \epsilon'.$$

Then

$$\frac{a_i}{N + C_{t-1}(\lambda,p,r,\epsilon')/\alpha} \leq \frac{|Y_i|}{\sum_{i=0}^k |Y_i|} \leq \frac{a_i + C_{t-1}(\lambda,p,r,\epsilon')}{N}.$$

Thus

$$\left| a_i - \frac{|Y_i|}{\sum_{i=0}^k |Y_i|}\right| < \frac{C_{t-1}(\lambda,p,r,\epsilon')}{\alpha N} \leq \epsilon'.$$

Let G be the disjoint union of the Y_i. Hence it holds, according to Lemma 3.23

$$\|\mathbf{F} - G\|_{p,r}^{\mathrm{local}} \leq 2c_{r,p}\epsilon' < \epsilon.$$

Moreover, $N \leq |G| \leq N + C_{t-1}(\lambda,p,r,\epsilon')/\alpha \leq N + C_t(\lambda,p,r,\epsilon)$.

If $a_0 < \epsilon'$ we consider $Y_1,\dots,Y_k$ as above, but Y_0 is chosen with the only conditions that $|Y_0| \leq C_0(\lambda,p,r,\epsilon') \leq C_{t-1}(\lambda,p,r,\epsilon')$ and $Y_0 \equiv^r \mathbf{L}_0$. (Actually, Y_0 can be chosen so that $|Y_0|$ is bounded by a function of λ, p, and r only.) Let G be the disjoint union of the Y_i. Let $\hat{\mathbf{L}}_\mathbf{0}$ be the modeling with relational sample space $\mathbf{L}_0$ and probablity measure $\nu_{\hat{\mathbf{L}}_\mathbf{0}} = a_0^{-1}\nu_{\mathbf{L}}|L_0$ if $a_0 > 0$, and any probability measure if $a_0 = 0$ (for instance the discrete probability measure concentrated on $r_{\hat{\mathbf{L}}_\mathbf{0}}$). Let $\mathbf{F} \in \mathcal{F}^{(h)}$ be the forest defined from $\mathbf{F} = \coprod_{i=0}^k(\hat{\mathbf{L}}_\mathbf{i}, a_i)$ by making the component $\hat{\mathbf{L}}_\mathbf{0}$ special. It is clear that $\mathbf{L} = \mathsf{I}_{F\to Y}(\mathbf{F})$. Then, according to Lemma 3.23

$$\begin{aligned}\|\mathbf{F} - G\|_{p,r}^{\mathrm{local}} &\leq c_{r,p}\big(\epsilon' + \sum_{i=1}^k a_i\|\hat{\mathbf{L}}_\mathbf{i} - Y_i\|_{p,r}^{\mathrm{local}} + a_0\big)\\ &< c_{r,p}\big(2\epsilon' + \sup_{1\leq i\leq k}\|\hat{\mathbf{L}}_\mathbf{i} - Y_i\|_{p,r}^{\mathrm{local}}\big)\\ &\leq 3c_{r,p}\epsilon' = \epsilon.\end{aligned}$$

and, as above, $N \leq |G| \leq N + C_t(\lambda,p,r,\epsilon)$. Now, let $Y_\epsilon = \mathsf{I}_{F\to Y}(G)$. As $\mathsf{I}_{F\to Y}$ is basic and quantifier-free, and as $\mathsf{I}_{F\to}(\mathbf{Y}) = \mathbf{L}$ it holds $\|\mathbf{L} - Y_\epsilon\|_{p,r}^{\mathrm{local}} < \epsilon$ and $N \leq |Y| \leq N + C_t(\lambda,p,r,\epsilon)$. □

THEOREM 4.35. *A modeling* $\mathbf{L}$ *is the* FO*-limit of an* FO*-convergent sequence* $(Y_n)_{n\in\mathbb{N}}$ *of finite colored rooted trees with height at most* h *if and only if*

- L *is a colored rooted tree with height at most* h*,*
- L *satisfies the FMTP.*

4.3. Limit of Graphs with Bounded Tree-depth

Let Y be a rooted forest. The vertex x is an *ancestor* of y in Y if x belongs to the path linking y and the root of the tree of Y to which y belongs to. The *closure* Clos(Y) of a rooted forest Y is the graph with vertex set $V(Y)$ and edge set $\{\{x,y\} : x \text{ is an ancestor of } y \text{ in } Y, x \neq y\}$. The *height* of a rooted forest is the maximum number of vertices in a path having a root as an extremity. The *tree-depth* $\mathrm{td}(G)$ of a graph G is the minimum height of a rooted forest Y such that $G \subseteq \mathrm{Clos(Y)}$. This notion is defined in [59] and studied in detail in [71]. In particular, graphs with bounded tree-depth serve as building blocks for *low tree-depth decompositions*, see [60, 61, 62]. It is easily checked that for each integer t the property $\mathrm{td}(G) \leq t$ is first-order definable. It follows that for each integer t there exists a first-order formula ξ with a single free variable such that for every graph G and every vertex $v \in G$ it holds:

$$G \models \xi(v) \quad\iff\quad \mathrm{td}(G) \leq t \text{ and } \mathrm{td}(G - v) < \mathrm{td(G)}.$$

Let $t \in \mathbb{N}$. We define the basic interpretation scheme I_t, which interprets the class of connected graphs with tree-depth at most t in the class of 2^{t-1}-colored rooted trees: given a 2^{t-1}-colored rooted tree $\mathbf{Y}$ (where colors are coded by $t-1$ unary relations $C_1, \dots, C_{t-1}$), the vertices $u, v \in Y$ are adjacent in $\mathsf{I}_t(\mathbf{Y})$ if the there is an integer i in $1, \dots, t-1$ such that $\mathbf{Y} \models C_i(v)$ and u is the ancestor of v at height i or $\mathbf{Y} \models C_i(u)$ and v is the ancestor of u at height i.

THEOREM 4.36. *Let $(\mathbf{G}_n)_{n\in\mathbb{N}}$ be an* FO*-convergent sequence of finite colored graphs with tree-depth at most h. Then there exists a colored rooted tree modeling $\mathbf{L} \in \mathcal{Y}^{(h)}$ satisfying the FMTP, such that the modeling $\mathbf{G} = \mathsf{I}_h(\mathbf{L})$ has tree-depth at most h and is a modeling* FO*-limit of the sequence $(\mathbf{G}_n)_{n\in\mathbb{N}}$.*

Conversely, if there is colored rooted tree modeling $\mathbf{L} \in \mathcal{Y}^{(h)}$ satisfying the FMTP and if $\mathbf{G} = \mathsf{I}_h(\mathbf{L})$, then there is an FO*-convergent sequence $(\mathbf{G}_n)_{n\in\mathbb{N}}$ of finite colored graphs with tree-depth at most h, such that $\mathbf{G}$ is a modeling* FO*-limit of $(\mathbf{G}_n)_{n\in\mathbb{N}}$.*

PROOF. For each $\mathbf{G}_n$, there is a colored rooted tree $\mathbf{Y}_n \in \mathcal{Y}^{(h)}$ such that $\mathbf{G}_n = \mathsf{I}_h(\mathbf{Y}_n)$. By compactness, the sequence $(\mathbf{Y}_n)_{n\in\mathbb{N}}$ has a converging subsequence $(\mathbf{Y}_{i_n})_{n\in\mathbb{N}}$, which admits a modeling FO-limit $\mathbf{Y}$ (according to Theorem 4.32), and it follows from Lemma 3.11 that $\mathsf{I}_h(\mathbf{Y})$ is a modeling FO-limit (with tree-depth at most h) of the sequence $(\mathbf{G}_{i_n})_{n\in\mathbb{N}}$, hence a modeling FO-limit of the sequence $(\mathbf{G}_n)_{n\in\mathbb{N}}$. □

CHAPTER 5

Concluding Remarks

5.1. Selected Problems

We hope that the theory developed here will encourage further researches. Here we list a sample of related problems

The first problem concern existence of modeling FO-limits. Recall that a class $\mathcal{C}$ is *nowhere dense* [65, 66, 67, 69] if, for every integer d there is an integer N such that the d-subdivision of K_N is not a subgraph of a graph in $\mathcal{C}$. We have proven, see Theorem 3.32, that if a monotone class $\mathcal{C}$ is such that every FO-convergent sequence of graphs in $\mathcal{C}$ has a modeling FO-limit, then $\mathcal{C}$ is nowhere dense. It is thus natural to ask whether the converse statement holds.

Problem 5.1. Let $\mathcal{C}$ be a nowhere dense class of graphs. Is it true that every FO-convergent sequence $(G_n)_{n\in\mathbb{N}}$ of finite graphs in $\mathcal{C}$ admit a modeling FO-limit?

Aldous-Lyons conjecture [5] states that every unimodular distribution on rooted countable graphs with bouded degree is the limit of a bounded degree graph sequence. One of the reformulations of this conjecture is that every graphing is an $\mathrm{FO}^{\mathrm{local}}$ limit of a sequence of finite graphs. The importance of this conjecture appears, for instance, in the fact that it would imply that all groups are sofic, which would prove a number of famous conjectures which are proved for sofic groups but still open for all groups.

The next problem is a possible strengthening of the conjecture.

Problem 5.2. Is every graphing $\mathbf{G}$ with the finite model property an FO-limit of a sequence of finite graphs?

Although the existence of a modeling FO-limit for FO-convergent sequences of graphs with bounded tree-depth follows easily from our study of FO-convergent sequence of rooted colored trees, the inverse theorem is more difficult. Indeed, if we would like to extend the inverse theorem for rooted colored trees to bounded tree-depth modelings, we naturally have to address the following question:

Problem 5.3. Is it true that there is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every graph modeling $\mathbf{L}$ with tree-depth at most t there exists a rooted colored tree modeling $\mathbf{Y}$ with height at most $f(t)$ such that $\mathbf{L} = \mathsf{I}_{f(t)}(\mathbf{Y})$, where the I_h (for $h \in \mathbb{N}$) are the basic intepretation schemes introduced in Section 4.3?

5.1.1. Classes with Bounded SC-depth.

We can generalize our main construction of limits to other tree-like classes. For example, in a similar way that we obtained a modeling FO-limit for FO-convergent sequences of graphs with bounded tree-depth, it is possible to get a modeling FO-limit for FO-convergent sequences of graphs with bounded SC-depth, where SC-depth is defined as follows [36]:

Let G be a graph and let $X \subseteq V(G)$. We denote by $\overline{G}^X$ the graph G' with vertex set $V(G)$ where $x \neq y$ are adjacent in G' if (i) either $\{x, y\} \in E(G)$ and $\{x, y\} \not\subseteq X$, or (ii) $\{x, y\} \notin E(G)$ and $\{x, y\} \subseteq X$. In other words, $\overline{G}^X$ is the graph obtained from G by complementing the edges on X.

Definition 5.1 (SC-depth). We define inductively the class $\mathcal{SC}(n)$ as follows:

- We let $\mathcal{SC}(0) = \{K_1\}$;
- if $G_1, \ldots, G_p \in \mathcal{SC}(n)$ and $H = G_1 \dot\cup \ldots \dot\cup G_p$ denotes the disjoint union of the G_i, then for every subset X of vertices of H we have $\overline{H}^X \in \mathcal{SC}(n+1)$.

The *SC-depth* of G is the minimum integer n such that $G \in \mathcal{SC}(n)$.

5.1.2. Classes with Bounded Expansion. A graph H is a *shallow topological minor* of a graph G at depth t if some $\leq 2t$-subdivision of H is a subgraph of G. For a class $\mathcal{C}$ of graphs we denote by $\mathcal{C} \,\widetilde{\nabla}\, t$ the class of all shallow topological minors at depth t of graphs in $\mathcal{C}$. The class $\mathcal{C}$ has *bounded expansion* if, for each $t \geq 0$, the average degrees of the graphs in the class $\mathcal{C} \,\widetilde{\nabla}\, t$ is bounded, that is (denoting $\overline{\mathrm{d}}(G)$ the average degree of a graph G):

$$(\forall t \geq 0) \qquad \sup_{G \in \mathcal{C} \,\widetilde{\nabla}\, t} \overline{\mathrm{d}}(G) < \infty.$$

The notion of classes with bounded expansion were introduced by the authors in [57, 58, 60], and their properties further studied in [61, 62, 23, 24, 63, 64, 66, 67, 71, 72] and in the monograph [68]. Particularly, classes with bounded expansion include classes excluding a topological minor, like classes with bounded maximum degree, planar graphs, proper minor closed classes, etc.

Classes with bounded expansion have the characteristic property that they admit special decompositions — the so-called *low tree-depth decompositions* — related to tree-depth:

Theorem 5.2 ([58, 60]). *Let $\mathcal{C}$ be a class of graph. Then $\mathcal{C}$ has bounded expansion if and only if for every integer $p \in \mathbb{N}$ there exists $N(p) \in \mathbb{N}$ such that the vertex set of every graph $G \in \mathcal{C}$ can be partitioned into at most $N(p)$ parts in such a way that the subgraph of G induced by any $i \leq p$ parts has tree-depth at most i.*

This decomposition theorem is the core of linear-time first-order model checking algorithm proposed by Dvořák, Kráľ, and Thomas [25, 26]. In their survey on methods for algorithmic meta-theorems [37], Grohe and Kreutzer proved that (in a class with bounded expansion) it is possible eliminate a universal quantification by means of the additions of a bounded number of new relations while preserving the Gaifman graph of the structure.

By an inductive argument, we deduce that for every integer p, r and every class $\mathcal{C}$ of λ-structure with bounded expansion, there is a signature $\lambda^+ \supseteq \lambda$, such that every λ-structure $\mathbf{A} \in \mathcal{C}$ can be lifted into a λ^+-structure $\mathbf{A}^+$ with same Gaifman graph, in such a way that for every first-order formula $\phi \in \mathrm{FO}_p(\lambda)$ with quantifier rank at most r there is an existential formula $\widetilde{\phi} \in \mathrm{FO}_p(\lambda^+)$ such that for every $v_1, \ldots, v_p \in A$ it holds

$$\mathbf{A} \models \phi(v_1, \ldots, v_p) \quad \iff \quad \mathbf{A}^+ \models \widetilde{\phi}(v_1, \ldots, v_p).$$

Moreover, by considering a slightly stronger notion of lift if necessary, we can assume that $\widetilde{\phi}$ is a local formula. We deduce that there is an integer $q = q(\mathcal{C}, p, r)$ such that checking $\phi(v_1,\dots,v_p)$ can be done by considering satisfaction of $\widetilde{\psi}(v_1,\dots,v_p)$ in subgraphs induced by q color classes of a bounded coloration. Using a low-tree depth decomposition (and putting the corresponding colors in the signature λ^+), we get that there exists finitely many induced substructures $\mathbf{A}_I^+$ ($I \in \binom{[N]}{q}$) with tree-depth at most q and the property that for every first-order formula $\phi \in \mathrm{FO}_p(\lambda)$ with quantifier rank at most r there is an existential formula $\widetilde{\phi} \in \mathrm{FO}_p(\lambda^+)$ such that for every $v_1,\dots,v_p \in A$ with set of colors $I_0 \subseteq I$ it holds

$$\mathbf{A} \models \phi(v_1,\dots,v_p) \iff \exists I \in \binom{[N]}{q-p}:\ \mathbf{A}^+_{I\cup I_0} \models \widetilde{\phi}(v_1,\dots,v_p).$$

Moreover, the Stone pairing $\langle \phi, \mathbf{A}\rangle$ can be computed by inclusion/exclusion from stone pairings $\langle \phi, \mathbf{A}_I^+\rangle$ for $I \in \binom{[N]}{\leq q}$.

Thus, if we consider an FO converging-sequence $(\mathbf{A}_n)_{n\in\mathbb{N}}$, the tuple of limits of the λ^+-structures $(\mathbf{A}_n)_I^+$ behaves as a kind of approximation of the limit of the λ-structures $\mathbf{A}_n$. We believe that this presents a road map for considering more general limits of sparse graphs.

Acknowledgements

The authors would like to thank Pierre Charbit for suggesting to use a proof by induction for Lemma 4.4 and Cameron Freer for his help in proving that the FO-limit of a sequence of random graphs cannot be a modeling.

Bibliography

[1] S. Adams, *Trees and amenable equivalence relations*, Ergodic Theory Dynam. Systems **10** (1990), 1–14.

[2] H. Adler and I. Adler, *Interpreting nowhere dense graph classes as a classical notion of model theory*, European J. Combin. (2013), to appear.

[3] D. Aldous, *Representations for partially exchangeable arrays of random variables*, J. Multivar. Anal. **11** (1981), 581–598.

[4] ———, *Exchangeability and related topics*, École d'été de probabilités de Saint-Flour, XIII — 1983, Lecture Notes in Mathematics, vol. 1117, Springer, 1985, pp. 1–198.

[5] D. Aldous and R. Lyons, *Processes on unimodular random networks*, Electronic Journal of Probability **12** (2007), 1454–1508.

[6] T. Austin, *On exchangeable random variables and the statistics of large graphs and hypergraphs*, Probability Surveys **5** (2008), 80–145.

[7] I. Benjamini, H. Finucane, and R. Tessera, *On the scaling limit of finite vertex transitive graphs with large diameter*, arXiv:1203.5624, 2012.

[8] I. Benjamini and O. Schramm, *Recurrence of distibutional limits of finite planar graphs*, Electron. J. Probab. **6** (2001), no. 23, 13pp.

[9] A. Blass, G. Exoo, and F. Harary, *Paley graphs satisfy all first-order adjacency axioms*, J. Graph Theory **5** (1981), 435–439.

[10] B. Bollobás and A. Thomason, *Graphs which contain all small graphs*, European J. Combin. **2** (1981), 13–15.

[11] C. Borgs, J. Chayes, L. Lovász, V.T. Sós, and K. Vesztergombi, *Counting graph homomorphisms*, Topics in Discrete Mathematics (M. Klazar, J. Kratochvíl, M. Loebl, J. Matoušek, R. Thomas, and P. Valtr, eds.), Algorithms and Combinatorics, vol. 26, Springer Verlag, 2006, pp. 315–371.

[12] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, B. Szegedy, and K. Vesztergombi, *Graph limits and parameter testing*, Proc. 38th Annual ACM Symp. Principles of Dist. Comp., 2005, pp. 51–59.

[13] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztergombi, *Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing*, Adv. Math. **219** (2008), no. 6, 1801–1851.

[14] C. Borgs, J.T. Chayes, L Lovász, V.T. Sós, and K. Vesztergombi, *Convergent sequences of dense graphs II: Multiway cuts and statistical physics*, Annals of Mathematics **176** (2012), 151–219.

[15] P.J. Cameron and J. Nešetřil, *Homomorphism-homogeneous relational structures*, Combin. Probab. Comput. **15** (2006), 91–103.

[16] A.K. Chandra and P.M. Merlin, *Optimal implementation of conjunctive queries in relational databases*, Proc. 9th ACM Symp. on Theory of Computing, 1977, pp. 77–90.

[17] P. Charbit, L. Hosseini, J. Nešetřil, and P. Ossona de Mendez, *On m-partite cographs*, in preparation, 2013.

[18] G. Cherlin, *Two problems on homogeneous structures, revisited*, Model Theoretic Methods in Finite Combinatorics (M. Grohe and J.A. Makowsky, eds.), vol. 558, American Mathematical Society, 2011, pp. 319–415.

[19] F. R. K. Chung, R. L. Graham, and R. M. Wilson, *Quasi-random graphs*, Combinatorica **9** (1989), no. 4, 345–362.

[20] P. Diaconis, S. Holmes, and S. Janson, *Threshold graph limits and random threshold graphs*, Internet Mathematics **5** (2009), no. 3, 267–318.

[21] P. Diaconis and S. Janson, *Graph limits and exchangeable random graphs.*, Rend. Mat. Appl., VII. Ser. **28** (2008), no. 1, 33–61 (English).

[22] R. Diestel, *Graph theory*, Springer Verlag, 1997.

[23] Z. Dvořák, *Asymptotical structure of combinatorial objects*, Ph.D. thesis, Charles University, Faculty of Mathematics and Physics, 2007.

[24] ———, *On forbidden subdivision characterizations of graph classes*, European J. Combin. **29** (2008), no. 5, 1321–1332.

[25] Z. Dvořák and D. Král', *Algorithms for classes of graphs with bounded expansion*, Lecture Notes in Computer Science **5911 LNCS** (2010), 17–32.

[26] Z. Dvořák, D. Král', and R. Thomas, *Deciding first-order properties for sparse graphs*, 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2010), 2010, pp. 133–142.

[27] A. Ehrenfeucht, *An application of games to the completeness problem for formalized theories*, Fund. Math. **49** (1961), 129–141.

[28] G. Elek, *Note on limits of finite graphs*, Combinatorica **27** (2007), 503–507.

[29] G. Elek and B. Szegedy, *Limits of hypergraphs, removal and regularity lemmas. A non-standard approach*, arXiv:0705.2179v1 [math.CO], 2007.

[30] P. Erdős and A. Rényi, *The evolution of random graphs*, Magyar Tud. Akad. Mat. Kutató Int. Közl **5** (1960), 17–61.

[31] P. Erdős and A. Rényi, *Asymmetric graphs*, Acta Math. Acad. Sci. Hungar. **14** (1963), 295–315.

[32] R. Fraïssé, *Sur une nouvelle classification des systèmes de relations*, Comptes Rendus **230** (1950), 1022–1024.

[33] ———, *Sur quelques classifications des systèmes de relations*, Ph.D. thesis, University of Paris, 1953, published in Publications Scientifiques de l'Université d'Alger, series A 1 (1954), 35182., pp. 1022–1024.

[34] D. Gaboriau, *Invariant percolation and harmonic Dirichlet functions*, Geometric And Functional Analysis **15** (2005), 1004–1051.

[35] H. Gaifman, *On local and non-local properties*, Proceedings of the Herbrand Symposium, Logic Colloquium '81, 1982.

[36] R. Ganian, P. Hliněný, J. Nešetřil, J. Obdržálek, P. Ossona de Mendez, and R. Ramadurai, *When trees grow low: Shrubs and fast* MSO_1, MFCS 2012, Lecture Notes in Computer Science, vol. 7464, Springer-Verlag, 2012, pp. 419–430.

[37] M. Grohe and S. Kreutzer, *Methods for algorithmic meta theorems*, Model Theoretic Methods in Finite Combinatorics, Contemporary mathematics, 2011, pp. 181–206.

[38] W. Hanf, *Model-theoretic methods in the study of elementary logic*, The theory of models (J.W. Addison et al., ed.), North-Holland, 1965, pp. 132–145.

[39] H. Hatami and S. Norine, *The entropy of random-free graphons and properties*, arXiv:1205.3529, May 2012.

[40] P. Hell and J. Nešetřil, *Graphs and homomorphisms*, Oxford Lecture Series in Mathematics and its Applications, vol. 28, Oxford University Press, 2004.

[41] W. Hodges, *Model theory*, Cambridge University Press, 1993.

[42] D. Hoover, *Relations on probability spaces and arrays of random variables*, Tech. report, Institute for Advanced Study, Princeton, NJ, 1979.

[43] S. Janson, *Graph limits and hereditary properties*, arXiv:1102.3571, Mar 2013.

[44] ______, *Graphons, cut norm and distance, couplings and rearrangements*, New York Journal of Mathematics Monographs, vol. 4, European Mathematical Information Service, 2013.

[45] O. Kallenberg, *Probabilistic symmetries and invariance principles*, Probability and Its Applications, Springer, 2005.

[46] A.S. Kechris, S. Solecki, and S. Todorcevic, *Borel chromatic numbers*, Advances in Mathematics **141** (1999), 1–44.

[47] A.H. Lachlan and R.E. Woodrow, *Countable ultrahomogeneous undirected graphs*, Trans. AMS **262** (1980), 51–94.

[48] D. Lascar, *La théorie des modèles en peu de maux*, Cassini, 2009.

[49] M.C. Laskowski, *Vapnik-chervonenkis classes of definable sets*, J. London Math. Soc **45** (1992), 377–384.

[50] P. Loeb, *Conversion from nonstandard to standard measure spaces and applications in probability theory*, Trans. AMS **211** (1975), 113–122.

[51] J. Łoś, *Quelques remarques, théorèmes et problèmes sur les classes définissables d'algèbres*, Mathematical interpretation of formal systems, Studies in logic and the foundations of mathematics, North-Holland, 1955.

[52] L Lovász, *Very large graphs*, arXiv:0902.0132, Feb. 2009.

[53] ______, *Large networks and graph limits*, Colloquium Publications, vol. 60, American Mathematical Society, 2012.

[54] L. Lovász and B. Szegedy, *Limits of dense graph sequences*, J. Combin. Theory Ser. B **96** (2006), 933–957.

[55] L Lovász and B. Szegedy, *Regularity partitions and the topology of graphons*, An irregular mind (Szemerédi is 70) (I. Bárány and J. Solymosi, eds.), Bolyai Society Mathematical Studies, vol. 21, Springer, 2010, pp. 415–446.

[56] R. Lyons, *Asymptotic enumeration of spanning trees*, Combinatorics, Probability and Computing **14** (2005), 491–522.

[57] J. Nešetřil and P. Ossona de Mendez, *The grad of a graph and classes with bounded expansion*, 7th International Colloquium on Graph Theory (A. Raspaud and O. Delmas, eds.), Electronic Notes in Discrete Mathematics, vol. 22, Elsevier, 2005, pp. 101–106.

[58] ______, *Linear time low tree-width partitions and algorithmic consequences*, STOC'06. Proceedings of the 38th Annual ACM Symposium on Theory of Computing, ACM Press, 2006, pp. 391–400.

[59] ______, *Tree depth, subgraph coloring and homomorphism bounds*, European Journal of Combinatorics **27** (2006), no. 6, 1022–1041.

[60] ______, *Grad and classes with bounded expansion I. decompositions*, European Journal of Combinatorics **29** (2008), no. 3, 760–776.

[61] ______, *Grad and classes with bounded expansion II. algorithmic aspects*, European Journal of Combinatorics **29** (2008), no. 3, 777–791.

[62] ———, *Grad and classes with bounded expansion III. restricted graph homomorphism dualities*, European Journal of Combinatorics **29** (2008), no. 4, 1012–1024.

[63] ———, *Fraternal augmentations, arrangeability and linearly bounded Ramsey numbers*, European Journal of Combinatorics **30** (2009), no. 7, 1696–1703.

[64] ———, *Extremal problems for sparse graphs*, An irregular mind (Szemerédi is 70) (I. Bárány and J. Solymosi, eds.), Bolyai Society Mathematical Studies, vol. 21, Springer, 2010, pp. 447–490.

[65] ———, *First order properties on nowhere dense structures*, The Journal of Symbolic Logic **75** (2010), no. 3, 868–887.

[66] ———, *From sparse graphs to nowhere dense structures: Decompositions, independence, dualities and limits*, European Congress of Mathematics, European Mathematical Society, 2010, pp. 135–165.

[67] ———, *Sparse combinatorial structures: Classification and applications*, Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (Hyderabad, India) (R. Bhatia and A. Pal, eds.), vol. IV, World Scientific, 2010, pp. 2502–2529.

[68] ———, *How many F's are there in G?*, European Journal of Combinatorics **32** (2011), no. 7, 1126–1141.

[69] ———, *On nowhere dense graphs*, European Journal of Combinatorics **32** (2011), no. 4, 600–617.

[70] ———, *A model theory approach to structural limits*, Commentationes Mathematicæ Universitatis Carolinæ **53** (2012), no. 4, 581–603.

[71] ———, *Sparsity (graphs, structures, and algorithms)*, Algorithms and Combinatorics, vol. 28, Springer, 2012, 465 pages.

[72] J. Nešetřil, P. Ossona de Mendez, and D.R. Wood, *Characterizations and examples of graph classes with bounded expansion*, European Journal of Combinatorics **33** (2012), no. 3, 350–373.

[73] O. Pikhurko, *An analytic approach to stability*, Discrete Math. **310** (2010), 2951–2964.

[74] M.H. Stone, *The theory of representations of Boolean algebras*, Transactions of the American Mathematical Society **40** (1936), 37–111.