

Some classical model theoretic aspects of bounded shrub-depth classes

Abhisekh Sankaran

Department of Computer Science and Technology,
University of Cambridge, U.K.

{abhisekh.sankaran}@cl.cam.ac.uk

Abstract

We consider classes of arbitrary (finite or infinite) graphs of bounded shrub-depth, specifically the class $\text{TM}_{r,p}(d)$ of p -labeled arbitrary graphs whose underlying unlabeled graphs have tree models of height d and r labels. We show that this class satisfies an extension of the classical Löwenheim-Skolem property into the finite and for MSO. This extension being a generalization of the small model property, we obtain that the graphs of $\text{TM}_{r,p}(d)$ are pseudo-finite. In addition, we obtain as consequences entirely new proofs of a number of known results concerning bounded shrub-depth classes (of finite graphs) and $\text{TM}_{r,p}(d)$. These include the small model property for MSO with elementary bounds, the classical compactness theorem from model theory over $\text{TM}_{r,p}(d)$, and the equivalence of MSO and FO over $\text{TM}_{r,p}(d)$ and hence over bounded shrub-depth classes. The proof for the last of these is via an adaptation of the proof of the classical Lindström’s theorem from model theory characterizing FO over arbitrary structures.

1. INTRODUCTION

Classical model theory is a subject that studies mathematical structures for their properties that can be expressed in a formal language such as first order logic (FO). The structures considered are arbitrary, so finite or infinite (but are typically infinite). Two of the earliest results in the subject, which form its pillars even today, are the Löwenheim-Skolem theorem and the compactness theorem [1]. The Löwenheim-Skolem (LS) theorem states that for every infinite structure \mathcal{A} over a countable vocabulary and every infinite cardinal λ , there is a structure of size λ that is FO equivalent to \mathcal{A} , that further is a substructure of \mathcal{A} if λ is less than the size of \mathcal{A} (the “downward” part of the theorem, denoted DLS), and a superstructure of \mathcal{A} otherwise (the “upward” part, denoted ULS). On the other hand, the compactness theorem states that if every finite subset of an arbitrary set of FO sentences has a model, then the entire set has a model. The importance of these results can be gauged from a classic result by Lindström [9] who showed that for a reasonable notion of an abstract logic, if such a logic either satisfies both the DLS and compactness theorems, or satisfies the full LS theorem (so both its parts), then it is no more expressive than FO.

Shrub-depth is a structural parameter of classes of finite graphs that was introduced in [6] to admit fixed parameter tractable (FPT) algorithms for a variety of interesting algorithmic problems, with an *elementary* parameter dependence. The notion parameterizes classes of dense graphs by the height of their defining tree models, where informally, a tree model t of a graph G is a rooted tree whose leaves are the vertices of G , and these are assigned labels from a finite set. The presence or absence of an edge between two vertices of G is determined by the labels of these vertices in t and the distance between them in t . Since its inception in [6], the notion has seen a lot of active research for not just its algorithmic properties, but also its structural and logical aspects [5, 7, 2]. For instance, graphs of shrub-depth at most d admit small models for MSO sentences (of sizes bounded by elementary functions of the lengths of the sentences), and admit the equivalence of MSO and FO in terms of their expressive powers [5].

In [10], a finitary analogue of the downward Löwenheim-Skolem property was introduced. A class of structures has this property denoted \mathcal{L} -EBSP in [10] for a logic \mathcal{L} such as FO or MSO, if any large structure in the class contains a small $\mathcal{L}[m]$ -equivalent substructure that belongs to class. Here $\mathcal{L}[m]$ -equivalent means that the two structures agree on all \mathcal{L} sentences of quantifier nesting depth (or rank) at most m , and “small” means of size bounded by a computable function of m . Bounded shrub-depth classes satisfy MSO-EBSP (which can be shown to be equivalent to the small submodel property for MSO), and it turns out going further, that a more general property holds of these classes, one called the *logical fractal* property [10, Section 6]. This property asserts “logical self-similarity” at all “scales” less than the size of a given structure, in analogy with the notion of fractals studied in the natural sciences. More specifically, for a strictly increasing function $f : \mathbb{N} \rightarrow \mathbb{N}$, called a “scale function” in [10], define the i^{th} scale of f as the set of all numbers in the interval $(f(i), f(i + 1)]$. Then an \mathcal{L} -fractal is a class \mathcal{C} of structures for which for each natural number m , there is a scale function f_m such that for each structure \mathcal{A} in \mathcal{C} , there is an $\mathcal{L}[m]$ -equivalent substructure \mathcal{B} of \mathcal{A} in \mathcal{C} , at each scale less than the scale of \mathcal{A} . It turns out that bounded shrub-depth classes are MSO-fractals with computable scale functions. The \mathcal{L} -fractal property can be seen as a finitary \mathcal{L} -DLS property that is closer to the classical DLS property than \mathcal{L} -EBSP.

In this paper, we generalize significantly the above result for bounded shrub-depth classes, that further yields us various notable consequences. Specifically, we consider a finitary analogue of the full LS property, namely both its downward and upward parts, that therefore generalizes the logical fractal property. For f as above, define for an infinite cardinal λ , the λ^{th} scale of f simply as the set $\{\lambda\}$. A class \mathcal{C} of arbitrary structures (so now finite and infinite structures) is said to satisfy the *extended \mathcal{L} -Löwenheim-Skolem (\mathcal{L} -ELS) property* if there exists a scale function f such that for any structure \mathcal{A} in \mathcal{C} , if \mathcal{A} is at scale η of f for some cardinal η (that is, the size of \mathcal{A} belongs to the scale η), then for all cardinals λ , the following hold:

1. (Downward \mathcal{L} -ELS): If λ is at most η , then (i) if λ is finite, then there is a substructure of \mathcal{A} , at scale λ that is $\mathcal{L}[\lambda]$ -equivalent to \mathcal{A} , and (ii) if λ is infinite, there is a substructure of \mathcal{A} , at scale λ that is \mathcal{L} -equivalent to \mathcal{A} , where \mathcal{L} -equivalent means indistinguishable with respect to all sentences of \mathcal{L} .
2. (Upward \mathcal{L} -ELS): If λ is at least η , then (i) if η is finite, then there is a superstructure of \mathcal{A} at scale λ that is $\mathcal{L}[\eta]$ -equivalent to \mathcal{A} , and (ii) if η is infinite, then there is a

superstructure of \mathcal{A} at scale λ that is \mathcal{L} -equivalent to \mathcal{A} .

We call f a “witness function” for the MSO-ELS property of \mathcal{C} . From the perspective of fractals, the \mathcal{L} -ELS property can be seen as asserting “decreasing \mathcal{L} -self-similarity” going downwards from the size of a given structure and “increasing \mathcal{L} -self-similarity” going upwards. It is not difficult to see that for any given m , by adjusting the scales suitably, specifically by merging all scales of f that are at most m , the \mathcal{L} -ELS property, restricted to the finite, implies the \mathcal{L} -fractal property since $\mathcal{L}[m']$ -equivalence implies $\mathcal{L}[m]$ -equivalence for all $m' \geq m$. We now consider the extension of the notion of bounded shrub-depth to arbitrary graphs as introduced in [2], obtained by considering tree models of bounded height but arbitrary branching. Specifically, for $r, p, d \in \mathbb{N}$, we consider the class $\text{TM}_{r,p}(d)$ of arbitrary p -labeled graphs (that is, graphs whose vertices are labeled with labels from $\{1, \dots, p\}$) that admit arbitrary tree models of height d and r labels, and show as the central result of our paper, the following theorem.

Theorem 1.1. *The class $\text{TM}_{r,p}(d)$ satisfies the MSO-ELS property with elementary $((d+1)$ -fold exponential) witness functions.*

We prove the above theorem in two parts, the first where one of λ or η in the MSO-ELS definition is finite (Theorem 4.1), and the second when both are infinite (Proposition 4.2). Both of these parts crucially build upon the result that $\text{TM}_{r,p}(d)$ admits small MSO[m]-equivalent submodels whose sizes are bounded by a d -fold exponential function of m (Theorem 1). Thus in the parlance of model theory, the graphs of $\text{TM}_{r,p}(d)$ are *pseudo-finite*, in that every FO (indeed also MSO) sentence true in a graph G of $\text{TM}_{r,p}(d)$, is also true in a finite graph of $\text{TM}_{r,p}(d)$, and further even in infinitely many finite induced subgraphs of G if G is infinite.

As consequences of Theorem 1.1, we obtain entirely new proofs for a variety of known results in the literature concerning bounded shrub-depth classes. Firstly, it follows that shrub-depth d classes of finite graphs admit small models for MSO of elementary sizes (shown in [5]); consequently, the index of the MSO[m] equivalence relation over any shrub-depth d class is elementary again. Secondly, we obtain a new proof of the classical compactness theorem from model theory mentioned at the outset, over the class $\text{TM}_{r,p}(d)$. That is, we show that for any set T of MSO sentences over the vocabulary of p -labeled graphs, if every finite subset of T has a model in $\text{TM}_{r,p}(d)$, then all of T has a model in $\text{TM}_{r,p}(d)$; and further, a countable such model (Theorem 5.1). It can be seen following [2] that $\text{TM}_{r,p}(d)$ can be axiomatized (over all structures) by a single FO sentence, so that $\text{TM}_{r,p}(d)$ is an *elementary* class in the classical model-theoretic language, where elementary here means “axiomatizable by a set of FO sentences”. Since most classical model theoretic results, and in particular compactness, hold over elementary classes, the FO compactness theorem for $\text{TM}_{r,p}(d)$ follows already from literature. However the standard proofs of the compactness theorem are either via Gödel’s completeness theorem for FO, or Henkin models, or Skolemization or ultraproducts [1]. Our approach is completely different from all of these and goes via using the pseudo-finiteness of the graphs of $\text{TM}_{r,p}(d)$ to construct a chain under induced subgraph of models of increasingly larger subsets of the given theory T , whose union is a countable model of T . We believe this technique might be of independent interest.

The reader might wonder why we have not mentioned our MSO compactness theorem for $\text{TM}_{r,p}(d)$ to be more general than the FO compactness theorem. The reason is that, as

it turns out, MSO is no more expressive than FO over $\text{TM}_{r,p}(d)$ (Theorem 6.1). This fact has been shown for $\text{TM}_{r,1}(d)$ (so basically the unlabeled versions of graphs in $\text{TM}_{r,p}(d)$) in [2], and for the class of finite graphs in $\text{TM}_{r,1}(d)$ in [5], but we provide a novel proof of this result (and further also for $\text{TM}_{r,p}(d)$ for all $p > 1$) by adapting the ideas from the proof of the classical Lindström’s theorem mentioned at the outset, which indeed asserts a “logical collapse” to FO for any abstract logic having the DLS and compactness properties. Since MSO does satisfy the conditions of the mentioned abstract logic, it is natural to think that Lindström’s theorem can be employed as is, to prove the MSO-FO equivalence over $\text{TM}_{r,p}(d)$. But it turns out that cannot be done since the proof of Lindström’s theorem requires a compactness theorem for structures over the given vocabulary expanded with binary relation and function symbols. Therefore, we directly adapt and implement the ideas from Lindström’s proof, in our setting. We remark that our MSO compactness theorem, particularly the fact that it gives a *countable* model, turns out to be vital for this adaptation to go through.

The main tool at the heart of our results, is an “unbounded” version of Feferman-Vaught composition theorem proved by Elberfeld, Grohe and Tantau in [4]. The standard Feferman-Vaught composition theorem shows that evaluating an MSO sentence over the disjoint union of finitely many structures is equivalent to evaluating a finite set of MSO sentences over each of the individual structures, and then putting together the results of the evaluations using a Boolean propositional logic formula. In [4], this result is extended elegantly to infinitely many structures using FO instead of propositional logic. Specifically, for any MSO sentence Φ , to know whether it is true over the disjoint union of an arbitrary family \mathcal{F} of structures, it is equivalent to examine the truth of an FO sentence α_Φ over an MSO[m]-*type indicator* $\mathcal{I}_m(\mathcal{F})$ which is a structure over a monadic vocabulary, that contains all the information about the equivalence classes of the MSO[m] relation to which the structures in \mathcal{F} belong (indeed these classes constitute the mentioned vocabulary). The FO sentence α_Φ constructed inductively on the structure of Φ , can be seen as doing both the job of evaluating suitable MSO[m] sentences over the individual structures of \mathcal{F} and the job of compiling together the results of the individual evaluations. The key observation for our results is the fact that the vocabulary of the type indicator is *monadic*. Structures over monadic vocabularies are particularly nice when it comes to dealing with them using FO sentences of bounded rank, since, when they are sufficiently large, one can always shrink them (upto a certain threshold depending on the rank) or expand them (without any bounds, even infinitely) by respectively deleting or adding one element at a time without changing their FO[q] theory, where q is the considered rank. This simple fact is exploited over and over again to prove all our results mentioned above. Specifically, the results are first proved for trees of bounded height noting that trees are after all constructed inductively from forests of lesser height, and the latter lend themselves to using the mentioned composition theorem. Subsequently, these results for trees are transferred to $\text{TM}_{r,p}(d)$ using the FO interpretability of the latter in the former.

The organization of the paper is as follows. In Section 2, we provide the background and notation for the paper, also recall the Feferman-Vaught composition theorem from [4]. In Section 3, we show the small model property for MSO over $\text{TM}_{r,p}(d)$ with elementary bounds. In Section 4, we show the MSO-ELS property for $\text{TM}_{r,p}(d)$. In Section 5 we prove the MSO compactness theorem over $\text{TM}_{r,p}(d)$ and utilize this in Section 6 to show the MSO-FO equivalence over $\text{TM}_{r,p}(d)$. We present our conclusions in Section 7.

2. BACKGROUND

We assume the reader is familiar with the terminology and notation in connection with MSO. A vocabulary, typically denoted τ or σ , is a finite set of relation symbols. We denote the class of all MSO formulae over τ as $\text{MSO}(\tau)$ and its subclass of FO formulae as $\text{FO}(\tau)$. A sequence x_1, \dots, x_n of FO variables is denoted \bar{x} . An MSO formula φ whose free variables are in \bar{x} is denoted $\varphi(\bar{x})$. A sentence is a formula without free variables. The *quantifier rank*, or simply *rank*, of an MSO formula φ , denoted $\text{rank}(\varphi)$, is the maximum number of quantifiers (both first order and second order) appearing in any root to leaf path in the parse tree of the formula. We denote by $\text{MSO}[m]$ the class of all MSO formulae of rank at most m .

A σ -structure \mathcal{A} consists of a universe A equipped with interpretations of the predicates of σ . We denote by $\mathcal{B} \subseteq \mathcal{A}$ that \mathcal{B} is a substructure of \mathcal{A} , and by $\mathcal{A} \cong \mathcal{B}$ that \mathcal{A} and \mathcal{B} are isomorphic. We denote by \mathbb{N} the set of all natural numbers, by \mathbb{N}_+ the set $\mathbb{N} \setminus \{0\}$, and by $[p]$ the set $\{1, \dots, p\}$ for any $p \in \mathbb{N}_+$. For $p \in \mathbb{N}_+$, let $\tau = \sigma \cup \{P_1, \dots, P_p\}$ be a vocabulary obtained by expanding σ with p new unary relation symbols P_1, \dots, P_p . A *p-labeled σ -structure* is a τ -structure \mathcal{A} in which for every element a in A , there exists a unique $i \in [p]$ such that $a \in P_i^{\mathcal{A}}$. (It is allowed for $P_j^{\mathcal{A}}$ to be empty for some $j \in [p]$.) We say of the mentioned element a , that it is labeled with the label i . We consider simple, undirected and loop-free graphs in this paper, where such a graph G is a σ -structure for $\sigma = \{E\}$, in which the binary relation symbol E is interpreted as an irreflexive and symmetric relation. The graphs can have *arbitrary* cardinality (i.e. whose universe could be finite or infinite). We denote by $V(G)$ and $E(G)$, the vertex and edge sets of G . As above, a *p-labeled graph* G is graph each of whose vertices is assigned a unique label from $[p]$. If G_1 and G_2 are two p -labeled graphs, and $G_1 \subseteq G_2$, then we say G_1 is a *labeled induced subgraph* of G_2 .

A *tree* is a connected graph that does not contain any cycles. A rooted tree \mathbf{t} is a $\{E, \text{root}\}$ -structure whose $\{E\}$ -reduct is a tree and in which the unary relation symbol root is interpreted as a set consisting of a single element called the *root* of \mathbf{t} , and denoted $\text{root}(\mathbf{t})$. We can now talk of a *p-labeled rooted tree* as a rooted tree in which each node (including the root) is assigned a unique label from $[p]$. We shall often call p -labeled rooted trees, as simply trees when p is clear from context. For a tree \mathbf{t} and a node v of it, the *subtree of \mathbf{t} rooted at v* , denoted \mathbf{t}_v , is the $\{E, \text{root}\}$ -structure such that (i) the $\{E\}$ -reduct of \mathbf{t}_v is the substructure of \mathbf{t} induced by the set of all nodes u of \mathbf{t} for which v lies on the path from u to $\text{root}(\mathbf{t})$, and (ii) root is interpreted in \mathbf{t}_v as the set $\{v\}$. We say a tree \mathbf{s} is a subtree of \mathbf{t} if it is the subtree of \mathbf{t} rooted at some node v of \mathbf{t} (note the slight abuse of the standard ‘subtree’ terminology, intended for later convenience). If v is a child of $\text{root}(\mathbf{t})$, then we say \mathbf{t}_v is a *child subtree* of $\text{root}(\mathbf{t})$ in \mathbf{t} . A p -labeled rooted forest \mathbf{f} is a disjoint union of p -labeled rooted trees. If $(\mathbf{t}_i)_{i \in I}$ for an index set I is the family of p -labeled rooted trees constituting \mathbf{f} , then we write $\mathbf{f} = \uplus_{i \in I} \mathbf{t}_i$ where \uplus denotes disjoint union. A subtree \mathbf{t}_2 of a (p -labeled rooted) tree \mathbf{t}_1 is said to be *leaf-hereditary* if the roots of \mathbf{t}_2 and \mathbf{t}_1 are the same, and every leaf of \mathbf{t}_2 is also a leaf of \mathbf{t}_1 . It follows that if \mathbf{f}_i is the forest of p -labeled rooted trees obtained by removing the root of \mathbf{t}_i for $i \in \{1, 2\}$, then for every tree \mathbf{s}_2 of \mathbf{f}_2 , there exists a tree \mathbf{s}_1 of \mathbf{f}_1 such that \mathbf{s}_2 is a leaf-hereditary subtree of \mathbf{s}_1 . The *height* of a tree \mathbf{t} is the maximum root to leaf distance in \mathbf{t} . The height of a node v of a tree \mathbf{t} is the height of the subtree \mathbf{t}_v . Singleton trees have height 0. We denote by $\mathcal{T}_{d,p}$ the class of all (arbitrary cardinality) p -labeled rooted trees of height at most d .

2.1. SHRUB-DEPTH

We recall the notion of tree models from [3] and state it in its extended version for arbitrary cardinality graphs. For $r, d \in \mathbb{N}$, a *tree model of r labels and height d* for a graph G is a pair (\mathbf{t}, S) where \mathbf{t} is an $(r + 1)$ -labeled rooted tree of height d and $S \subseteq [r]^2 \times [d]$ is a set called the *signature* of the tree model such that:

1. The length of every root to leaf path is exactly d .
2. The set $V(G)$ is exactly the set of leaves of \mathbf{t} .
3. Each leaf is assigned a unique label from $[r]$ and all internal nodes are labeled $r + 1$.
4. For any $i, j \in [r]$ and $l \in [d]$, it holds that $(i, j, l) \in S$ if and only if $(j, i, l) \in S$.
5. For vertices $u, v \in V(G)$, if i and j are the labels of u and v seen as leaves of \mathbf{t} , and the distance between u and v in \mathbf{t} is $2l$, then $\{u, v\} \in E(G)$ iff $(i, j, l) \in S$. Observe that the distance between u and v in \mathbf{t} is an even number as all root to leaf paths are of length d , and l is thus the distance between u (or v) and the least common ancestor of u and v .

The class of all (arbitrary) tree models of r labels and height d is denoted $\text{Tree}_r(d)$ and the class of all (arbitrary) graphs that have tree models in $\text{Tree}_r(d)$ is denoted $\text{TM}_r(d)$. A class \mathcal{C} of arbitrary graphs is said to have *shrub-depth d* if $\mathcal{C} \subseteq \text{TM}_r(d)$ for some $r \geq 1$, and $\mathcal{C} \not\subseteq \text{TM}_{r'}(d-1)$ for any $r' \geq 1$. We say \mathcal{C} has bounded shrub-depth if $\mathcal{C} \subseteq \text{TM}_r(d)$ for some $r, d \geq 1$. For the purposes of our results, we extend the notion of tree models recalled above in a simple way to handle classes of p -labeled graphs whose underlying unlabeled graphs form classes having bounded shrub-depth. Given $r, p, d \in \mathbb{N}, p \geq 1$, let \mathbf{s} be a tree whose internal nodes are labeled $r + 1$ and each of whose leaf nodes is labeled with a pair (i, j) where $i \in [r]$ and where $j \in [p]$. For a set $S \subseteq [r]^2 \times [d]$, we say $(\mathbf{s}, S) \in \text{Tree}_{r,p}(d)$ if the tree \mathbf{t} obtained from \mathbf{s} by replacing the pair labeling any leaf node with the first component of the pair (so replacing every leaf label of the form (i, j) with just i), is such that $(\mathbf{t}, S) \in \text{Tree}_r(d)$. It is easy to see that (\mathbf{s}, S) defines a p -labeled graph G whose underlying unlabeled graph is the graph of $\text{TM}_r(d)$ for which (\mathbf{t}, S) is a tree model, and such that for any node $u \in V(G)$, if (i, j) is the label of u seen as a leaf node of \mathbf{s} , then the label of u in G is j . We say that (\mathbf{s}, S) is a tree-model of G . Denote by $\text{TM}_{r,p}(d)$ the class of arbitrary p -labeled graphs G for which there is a tree model in $\text{Tree}_{r,p}(d)$. It is easy to see that $\text{TM}_{r,p}(d)$ is a hereditary class, i.e. it is closed under induced subgraphs.

From the definitions above, we see that for any $(\mathbf{s}, S) \in \text{TM}_{r,p}(d)$, the tree \mathbf{s} can be seen as belonging to $\mathcal{T}_{d,r \cdot p+1}$ via the bijective function $f : [r] \times [p] \rightarrow [r \cdot p]$ that maps the pair (i, j) labeling any leaf node of \mathbf{s} to the number $(i - 1) \cdot p + j$, with every internal node of \mathbf{s} getting the label $r \cdot p + 1$. We will therefore also use the notation $\mathbf{s} \in \mathcal{T}_{d,r \cdot p+1}$. It is easy to write an FO sentence $\Omega_{r,p,d}$ over the vocabulary of $\mathcal{T}_{d,r \cdot p+1}$ that says that for a given tree from $\mathcal{T}_{d,r \cdot p+1}$, every root to leaf path has length exactly d , every internal node is labeled with the label $r \cdot p + 1$, and no leaf node is labeled with the label $r \cdot p + 1$. Then any model $\mathbf{s} \in \mathcal{T}_{d,r \cdot p+1}$ of $\Omega_{r,p,d}$ when equipped with a signature $S \subseteq [r]^2 \times [p]$ gives a tree model $(\mathbf{s}, S) \in \text{Tree}_{r,p}(d)$, and conversely for any tree model $(\mathbf{s}, S) \in \text{Tree}_{r,p}(d)$, we have \mathbf{s} , seen as a tree of $\mathcal{T}_{d,r \cdot p+1}$, models $\Omega_{r,p,d}$. Observe that the rank of $\Omega_{r,p,d}$ is

$O(d)$. For every signature $S \subseteq [r]^2 \times [p]$, there exists a tuple $\Xi_{S,p}$ of FO formulas given by $\Xi_{S,p} = (\xi_{V,S}(x), \xi_{E,S}(x,y), (\xi_{P_j,S}(x))_{1 \leq j \leq p})$ that when evaluated on a tree \mathfrak{s} satisfying $\Omega_{r,p,d}$ produces the graph $G \in \text{TM}_{r,p}(d)$ of which $(\mathfrak{s}, S) \in \text{Tree}_{r,p}(d)$ is a tree model. Specifically, (i) the formula $\xi_{V,S}(x)$ says that x is a leaf node in \mathfrak{s} , (ii) the formula $\xi_{E,S}(x,y)$ says that x and y are leaves of \mathfrak{s} and for some $(i,j,l) \in S$, it holds that i and j are resp. the labels of x and y , and the distance between x and y in \mathfrak{t} is exactly $2l$, and (iii) the formula $\xi_{P_j,S}(x)$ says that x is a leaf node whose label corresponds to the pair (i,j) for some $i \in [r]$. We call $\Xi_{S,p}$ an FO *interpretation* of $\text{TM}_{r,p}(d)$ in $\text{Tree}_{r,p}(d)$. Thus $\Xi_{S,p}$ defines a function from $\text{Tree}_{r,p}(d)$ to $\text{TM}_{r,p}(d)$, which we also denote as $\Xi_{S,p}$; so for G and \mathfrak{s} as above $G = \Xi_{S,p}(\mathfrak{s})$. Then $\text{TM}_{r,p}(d) = \bigcup_{S \subseteq [r]^2 \times [p]} \{\Xi_{S,p}(\mathfrak{s}) \mid (\mathfrak{s}, S) \in \text{Tree}_{r,p}(d)\}$. We now mention some important observations about $\Xi_{S,p}$. Firstly, the rank of $\Xi_{S,p}$, defined as the maximum rank of the formulae appearing in it, is $O(d)$. Next, if $G = \Xi_{S,p}(\mathfrak{s})$, then for any MSO formula φ in the vocabulary of G , there is an MSO formula that we denote $\Xi_{S,p}(\varphi)$ in the vocabulary of \mathfrak{s} such that $G \models \varphi$ iff $\mathfrak{s} \models \Xi_{S,p}(\varphi)$. (This is a special case of a more general result called the fundamental theorem of interpretations.) As a consequence, if \mathfrak{s}_1 and \mathfrak{s}_2 are two models of $\Omega_{r,p,d}$ and $q = \text{rank}(\Xi_{S,p})$ for a signature S , then $\mathfrak{s}_1 \equiv_{m+q, \text{MSO}} \mathfrak{s}_2$ implies $\Xi_{S,p}(\mathfrak{s}_1) \equiv_{m, \text{MSO}} \Xi_{S,p}(\mathfrak{s}_2)$. Finally for $(\mathfrak{s}, S) \in \text{Tree}_{r,p}(d)$, if \mathfrak{s}' is a leaf-hereditary subtree of \mathfrak{s} , then $\Xi_{S,p}(\mathfrak{s}')$ is a labeled induced subgraph of $\Xi_{S,p}(\mathfrak{s})$.

2.2. FEFERMAN-VAUGHT COMPOSITION

Let \mathcal{L} be one of the logic FO or MSO. Given $m \in \mathbb{N}$ and a τ -structures \mathcal{A} and \mathcal{B} , we say that \mathcal{A} and \mathcal{B} are $\mathcal{L}[m]$ -equivalent, denoted $\mathcal{A} \equiv_{m, \mathcal{L}} \mathcal{B}$, if \mathcal{A} and \mathcal{B} agree on all \mathcal{L} sentences of rank at most m . The relation $\equiv_{m, \text{MSO}}$ is an equivalence relation of finite index. Given a class \mathcal{C} of τ -structures and $m \in \mathbb{N}$, we let $\Delta_{m, \mathcal{C}}^{\mathcal{L}}$ denote the set of all equivalence classes of the $\equiv_{m, \mathcal{L}}$ relation over \mathcal{C} . We denote by $\delta_{m, \mathcal{C}}^{\mathcal{L}}(\mathcal{A})$ the equivalence class of $\Delta_{m, \mathcal{C}}^{\mathcal{L}}$ that contains \mathcal{A} . For every $\delta \in \Delta_{m, \mathcal{C}}^{\mathcal{L}}$, there exists an $\mathcal{L}[m]$ sentence Θ_δ that defines δ over \mathcal{C} . Define the function $\text{tower} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ as: $\text{tower}(0, n) = n$ and $\text{tower}(d, n) = 2^{\text{tower}(d-1, n)}$. Let $\iota(m, \mathcal{C}, \mathcal{L})$ denote the (finite) index of the $\equiv_{m, \mathcal{L}}$ relation over \mathcal{C} ; so $\iota(m, \mathcal{C}, \mathcal{L}) = |\Delta_{m, \mathcal{C}}^{\mathcal{L}}|$. The relation $\equiv_{m, \mathcal{L}}$ has a characterization in terms of Ehrenfeucht-Fraïssé games. We point the reader [8, Chapters 3 & 7] for the details concerning these.

Let τ and \mathcal{C} be as above. Let $\mathcal{F} = (\mathcal{A}_i)_{i \in I}$ be a family of structures of \mathcal{C} with disjoint universes, indexed by an index set I of an arbitrary cardinality. Let $m \in \mathbb{N}$ and $\tau_{m, \mathcal{C}}$ be the relational vocabulary consisting of a distinct unary predicate symbol T for each class $T \in \Delta_{m, \mathcal{C}}^{\text{MSO}}$, and containing no other predicate symbols. Note that we are using the same symbol T to denote both an equivalence class of $\equiv_{m, \text{MSO}}$ over \mathcal{C} , as well as the unary predicate symbol corresponding to the class for the ease of understanding; whether we mean T as a class or a predicate symbol will be clear from context. The MSO[m]-*type indicator* for the family \mathcal{F} is now defined as a $\tau_{m, \mathcal{C}}$ -structure $\mathcal{I}_m(\mathcal{F})$ such that (i) the universe of $\mathcal{I}_m(\mathcal{F})$ is I , and (ii) for $T \in \Delta_{m, \mathcal{C}}^{\text{MSO}}$, the interpretation of T in $\mathcal{I}_m(\mathcal{F})$ is the set $\{i \in I \mid T = \delta_{m, \mathcal{C}}^{\text{MSO}}(\mathcal{A}_i)\}$. Observe that for each $i \in I$, there is exactly one predicate $T \in \tau_{m, \mathcal{C}}$ such that i is in the interpretation of T in $\mathcal{I}_m(\mathcal{F})$. We now have the following theorem from [4]. (This is the special case of $w = 0$ and $L = \text{MSO}$ in [4, Theorem 14].)

Theorem 2.1 (Theorem 14, [4]). *Let \mathcal{C} be a class of structures over a vocabulary τ . For every MSO sentence Φ over τ of rank m , there exists an FO sentence α_Φ over $\tau_{m, \mathcal{C}}$ such*

that if $\mathcal{F} = (\mathcal{A}_i)_{i \in I}$ is a family of structures of \mathcal{C} with disjoint universes, indexed by an index set I of an arbitrary cardinality, then the following holds:

$$\mathcal{I}_m(\mathcal{F}) \models \alpha_\Phi \quad \text{if, and only if,} \quad \bigcup_{i \in I} \mathcal{A}_i \models \Phi$$

Further, if $\tilde{\mathcal{C}}$ represents the class of structures of \mathcal{C} expanded with (all possible interpretations of) m new unary predicate symbols, then the rank of α_Φ is $O((\iota(m, \tilde{\mathcal{C}}, \text{MSO}))^{m+1})$.

Remark 2.2. In [4], the result is actually stated for Φ which does not contain any FO variables and whose atomic formulae, instead of being the usual atomic formulae (of the form $x_1 = x_2$, $Y(x)$ for an MSO variable Y , and $R(x_1, \dots, x_r)$ where $R \in \tau$ and x, x_1, \dots, x_r are FO variables), are instead of the “second order” forms $\text{empty}(X)$ and $\text{elem}(X_1, \dots, X_r, Z)$ where X, X_1, \dots, X_r are MSO variables and Z is either an MSO variable or a predicate of τ , with r being the arity of Z . The semantics for these atomic forms is intuitive: $\text{empty}(X)$ holds for set P if P is empty, and $\text{elem}(X_1, \dots, X_r, Z)$ holds for $X_i = P_i$ and $Z = Q$ if $|P_i| = 1$ for $1 \leq i \leq r$ and $P_i \times \dots \times P_r \subseteq Q$. Every “usual” MSO formula can be converted into an equivalent MSO formula over the mentioned second order atomic formulae, without any change of quantifier rank (see [4, page 4]). We have hence recalled [4, Theorem 14] in the form stated above in Theorem 2.1 which features Φ as a usual MSO formula, since these are the kinds of formulae which we work with in this paper.

We provide here the justification for the last statement of Theorem 2.1 which does not appear explicitly in [4] but is indeed a consequence of the proof of [4, Theorem 14]. We refer the reader to [4, Section 3.1, pp. 6 – 9] to find the formulae and other constructions we refer to in our description here.

We first observe in the proof of Lemma 8 of [4], that the “capping” constant C for the formula α is simply the rank of α , since α is an FO sentence over a monadic vocabulary τ_Σ (and we have also seen a similar such result in Lemma 3.3). Then the number n mentioned in the proof is at most $\text{rank}(\alpha) \cdot |\tau_\Sigma|$, whereby the rank of β_z , and hence the rank of β_α , is at most $\text{rank}(\alpha) \cdot |\text{vocab}(\alpha)| + 1$, where $\text{vocab}(\alpha)$ denotes the vocabulary of α , namely τ_Σ . Call this observation (*).

We now come to the proof of Theorem 14 of [4], and make the following observations about the rank of α_Φ following the inductive construction of α_Φ as given in the proof. For the base cases, since $\gamma_\Psi^L(i)$ is a quantifier-free formula for any Ψ , we get that if $\Phi := \text{empty}(X)$, then rank of α_Φ is 1, and if $\Phi := \text{elem}(X_1, \dots, X_r, Z)$ or $\Phi := \text{elem}(X_1, \dots, X_r, R)$, then the rank of α_Φ is 2 since the width w is 0 by our assumption. If Φ is a Boolean combination of a set of formulae, then $\text{rank}(\alpha_\Phi)$ is the maximum of the ranks of the formulae in the mentioned set. We now come to the non-trivial case when $\Phi := \exists X(\Phi')$.

We see from [4, page 9, para 2] that the rank of α_Φ is the maximum of the ranks of α_C where α_C is obtained from $\beta_{\alpha_{\Phi'}}$ (denoted as simply β in the proof as a shorthand) by substituting the atoms $T(i)$ with the quantifier-free formula $\gamma_\Psi^L(i)$ for a suitable Ψ . Then $\text{rank}(\alpha_\Phi) = \text{rank}(\beta_{\alpha_{\Phi'}})$. It follows from (*) above that $\text{rank}(\beta_{\alpha_{\Phi'}}) \leq \text{rank}(\alpha_{\Phi'}) \cdot |\text{vocab}(\alpha_{\Phi'})| + 1$; call this inequality (**). Let $\equiv_{r,L}$ denote the equivalence relation that relates two structures over the same vocabulary iff they agree on all L sentences (over the vocabulary of the structures) of rank at most r . Then the vocabulary of $\alpha_{\Phi'}$ is the set of all equivalence classes of the $\equiv_{q-1,L}$ relation over all structures over the vocabulary of the

family F (where F is as in the statement of [4, Theorem 14]), expanded with (all possible interpretations of) d set predicates, where d is the number of free variables of Φ' and $q-1$ is the rank of Φ' . Here we now importantly observe that if the structures of the family F come from a class \mathcal{C} , then it is sufficient to consider just those equivalence classes of the $\equiv_{q-1, L}$ relation that are non-empty when restricted to the structures of \mathcal{C} expanded with d set predicates. Then if Φ' is a subformula of a rank m L sentence Φ over a vocabulary τ and we are interested only in a given class \mathcal{C} of τ -structures and expansions of these with set predicates, then the size of the vocabulary of $\alpha_{\Phi'}$ is at most the index of the $\equiv_{m, L}$ equivalence relation over the class $\tilde{\mathcal{C}}$ of structures of \mathcal{C} expanded with (all possible interpretations of exactly) m set predicates. Applying this observation to (**) iteratively, we then get that if $L = \text{MSO}$ and $\lambda = \iota(m, \tilde{\mathcal{C}})$ then

$$\begin{aligned} \text{rank}(\alpha_{\Phi}) &\leq 1 + \lambda \cdot \left(1 + \lambda \cdot \left(1 + \dots \cdot (1 + 2 \cdot \lambda)\right)\right) \\ &\leq 1 + \lambda + \dots + \lambda^{m-1} + 2 \cdot \lambda^m \\ &\leq 2 \cdot \lambda^{m+1} \\ &= O((\iota(m, \tilde{\mathcal{C}}))^{m+1}) \end{aligned}$$

showing the last statement of Theorem 2.1.

3. SMALL MODELS AND ELEMENTARY BOUNDS

Following is the central result of this section.

Theorem 3.1. *Let $d, p, r \in \mathbb{N}$ be given. There exists an increasing function $h : \mathbb{N} \rightarrow \mathbb{N}$ such that the following are true for each $m \in \mathbb{N}_+$.*

1. *For every graph $G \in \text{TM}_{r,p}(d)$, there exists $H \in \text{TM}_{r,p}(d)$ such that (i) $H \subseteq G$. (ii) $|H|$ is at most $\text{tower}(d, h(d) \cdot m \cdot (m + \log r + \log p))$, and (ii) $H \equiv_{m, \text{MSO}} G$.*
2. *The index $\iota(m, \text{TM}_{r,p}(d), \text{MSO})$ of the $\equiv_{m, \text{MSO}}$ relation over $\text{TM}_{r,p}(d)$ is at most $\text{tower}(d+1, h(d) \cdot m^2 \cdot (\log r + \log p)^2)$.*

We prove Theorem 3.1 by first show the following core result for $\mathcal{T}_{d,p}$, and then transferring the latter to Theorem 3.1 using FO interpretations.

Theorem 3.2. *Let $d, p \in \mathbb{N}$ be given. There exists an increasing function $g : \mathbb{N} \rightarrow \mathbb{N}$ such that if $\zeta_{d,p} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ is the function given by $\zeta_{d,p}(n_1, n_2) = \text{tower}(n_2, g(d) \cdot (n_1 + 1) \cdot (n_1 + \log p))$, then the following are true for each $m \in \mathbb{N}_+$.*

1. *For every tree $\mathfrak{t} \in \mathcal{T}_{d,p}$, there exists a leaf-hereditary subtree \mathfrak{t}' of \mathfrak{t} such that (i) the heights of \mathfrak{t}' and \mathfrak{t} are the same, (ii) $|\mathfrak{t}'|$ is at most $\zeta_{d,p}(m, d)$, and (ii) $\mathfrak{t}' \equiv_{m, \text{MSO}} \mathfrak{t}$.*
2. *The index $\iota(m, \mathcal{T}_{d,p}, \text{MSO})$ of the $\equiv_{m, \text{MSO}}$ relation over $\mathcal{T}_{d,p}$ is at most $\zeta_{d,p}(m, d+1)$ if $d \geq 1$, and is p if $d = 0$.*

Proof of Theorem 3.1. Since $G \in \text{TM}_{r,p}(d)$, from Section 2 there exists a tree model $(\mathfrak{s}, S) \in \text{Tree}_{r,p}(d)$ for G and an FO interpretation $\Xi_{S,p}$ such that $G = \Xi_{S,p}(\mathfrak{s})$. Let $q = \max\{\max\{\text{rank}(\Xi_{S,p}) \mid S \subseteq [r]^2 \times [p]\}, \text{rank}(\Omega_{r,p,d})\}$ where $\Omega_{r,p,d}$ is as defined in

Section 2.1; we know that $q = O(d)$. Since s can be seen as a tree in $\mathcal{T}_{d,r \cdot p+1}$, by Theorem 3.2(1), there exists a leaf-hereditary subtree s' of s such that (i) $|s'| \leq \zeta_{d,r \cdot p+1}(q+m, d)$, and (ii) $s' \equiv_{q+m, \text{MSO}} s$. Then s' models $\Omega_{r,p,d}$ and hence $(s', S) \in \text{Tree}_{r,p}(d)$ is a tree model for $H = \Xi_{S,p}(s') \in \text{TM}_{r,p}(d)$. From the properties of the interpretation $\Xi_{S,p}$ (cf. Section 2.1), we infer the following: (i) Since s' is a leaf-hereditary subtree of s , we have $H \subseteq G$; (ii) Since $V(H)$ is the set of leaves of s' , we have $|H| \leq |s'| \leq \zeta_{d,r \cdot p+1}(q+m, d) = \text{tower}(d, g(d) \cdot (q+m+1) \cdot (q+m+\log(r \cdot p+1))) \leq \text{tower}(d, g(d) \cdot q^2 \cdot m \cdot (m+\log r+\log p)) \leq \text{tower}(d, h(d) \cdot m \cdot (m+\log r+\log p))$ where $h(d) = c_0 \cdot g(d) \cdot d^2$, $q \leq c_1 \cdot d$ and $c_0 \geq (c_1)^2$; (iii) Since $s' \equiv_{q+m, \text{MSO}} s$, we have $H \equiv_{m, \text{MSO}} G$.

We now look at the index of the $\equiv_{m, \text{MSO}}$ relation over $\text{TM}_{r,p}(d)$. For a given signature $S \subseteq [r]^2 \times [p]$, denote $\text{TM}_{r,p}^S(d)$ denote the subclass of $\text{TM}_{r,p}(d)$ of those graphs that have a tree model $(t, S) \in \text{Tree}_{r,p}(d)$, and $\text{Mod}(\Omega_{r,p,d})$ denote the class of models of $\Omega_{r,p,d}$ in $\mathcal{T}_{d,r \cdot p+1}$. Then $\Xi_{S,p}$ is a surjective map from $\text{Mod}(\Omega_{r,p,d})$ to $\text{TM}_{r,p}^S(d)$. For q as above, since any equivalence class of the $\equiv_{q+m, \text{MSO}}$ relation over $\text{Mod}(\Omega_{r,p,d})$ gets mapped by $\Xi_{S,p}$ to a subclass of an equivalence class of the $\equiv_{m, \text{MSO}}$ relation over $\text{TM}_{r,p}^S(d)$, we get by the surjectivity of Ξ_S that $i(m, \text{TM}_{r,p}^S(d), \text{MSO}) \leq i(q+m, \text{Mod}(\Omega_{r,p,d}), \text{MSO}) \leq i(c_1 \cdot d+m, \mathcal{T}_{d,r \cdot p+1}, \text{MSO})$. Then $i(m, \text{TM}_{r,p}(d), \text{MSO}) \leq \sum_{S \subseteq [r]^2 \times [p]} i(m, \text{TM}_{r,p}^S(d), \text{MSO}) \leq 2^{r^2 \cdot p} \cdot i(c_1 \cdot d+m, \mathcal{T}_{d,r \cdot p+1}, \text{MSO}) \leq 2^{r^2 \cdot p} \cdot \text{tower}(d+1, g(d) \cdot (c_1 \cdot d+m+1) \cdot (c_1 \cdot d+m+\log(r \cdot p+1))) \leq \text{tower}(d+1, h(d) \cdot (m \cdot (\log r + \log p))^2)$. \square

We now look at the proof of Theorem 3.2. The proof uses the following lemma in a central way. The lemma can be verified using a simple Ehrenfeucht-Fr ass e game argument.

Lemma 3.3. *Let σ be a finite vocabulary consisting of only monadic relation symbols. Let \mathcal{A} be an arbitrary (finite or infinite) σ -structure such that every element of \mathcal{A} is in the interpretation of exactly one predicate of σ . Let a be a given element of \mathcal{A} , and let $q \in \mathbf{N}_+$ and $q_1 = (q-1) \cdot |\sigma|$. Then the following are true:*

1. *If $|\mathcal{A}| > q_1$, then for every cardinal λ such that $q_1 < \lambda \leq |\mathcal{A}|$, there exists a substructure \mathcal{B} of \mathcal{A} such that (i) \mathcal{B} contains a , (ii) $|\mathcal{B}| = \lambda$, and (iii) $\mathcal{B} \equiv_{q, \text{FO}} \mathcal{A}$.*
2. *If $|\mathcal{A}| > q_1$, then for every cardinal $\lambda \geq |\mathcal{A}|$, there exists a σ -structure \mathcal{B} containing \mathcal{A} as substructure such that (i) $|\mathcal{B}| = \lambda$, and (ii) $\mathcal{B} \equiv_{q, \text{FO}} \mathcal{A}$. Further, for any $T \in \sigma$ such that $|T^{\mathcal{A}}| \geq q$, we can take \mathcal{B} to be a structure such that all elements of \mathcal{B} not in \mathcal{A} are in $T^{\mathcal{B}}$, and none of these elements is in the interpretation of any other predicate of σ in \mathcal{B} .*

We will also need the following lemma. The proof is easy and is skipped.

Lemma 3.4. *Let t_i for $i = 1, 2$ be a rooted tree and let f_i be the forest of rooted trees obtained by removing the root of t_i . Then $t_1 \equiv_{m, \text{MSO}} t_2$ if, and only if, $f_1 \equiv_{m, \text{MSO}} f_2$.*

Proof of Theorem 3.2. We prove the theorem by induction on d . The base case of $d = 0$ is trivial to see by taking $g(0) = 1$. Assume as induction hypothesis that the statement is true for $d-1$ for $d \geq 1$.

Consider a tree $t \in \mathcal{T}_{d,p}$ of height equal to d . Let f be the forest of p -labeled rooted trees obtained by removing the root of t . Specifically, let $f = \cup_{i \in I} s_i$ where $s_i \in \mathcal{T}_{d-1,p}$ and the (unrooted tree underlying) s_i is a child subtree of t , and I is a possibly infinite set. Let \mathcal{S}

be the class of rooted forests whose constituent trees belong to $\mathcal{T}_{d-1,p}$; so $f \in \mathcal{S}$. Consider now the MSO[m]-type indicator $\mathcal{I}_m(\mathcal{F})$ for the family $\mathcal{F} = (\mathbf{s}_i)_{i \in I}$ and the sentence $\Phi := \Theta_\delta$ for $\delta = \delta_{m,\mathcal{S}}^{\text{MSO}}(f)$ (the MSO[m] sentence axiomatizing the $\equiv_{m,\text{MSO}}$ equivalence class of f in \mathcal{S}). By Theorem 2.1, there exists an FO sentence α_Φ over the vocabulary $\tau_{m,\mathcal{T}_{d-1,p}}$ such that

$$\mathcal{I}_m(\mathcal{F}) \models \alpha_\Phi \quad \text{if, and only if,} \quad f \models \Phi$$

Now we know from Theorem 2.1 that for any MSO[m] sentence Ψ over the vocabulary of $\mathcal{T}_{d-1,p}$, the sentence α_Ψ given by the theorem has rank that is $O((\iota(m, \tilde{\mathcal{T}}_{d-1,p}, \text{MSO}))^{m+1})$ where $\tilde{\mathcal{T}}_{d-1,p}$ is the expansion of $\mathcal{T}_{d-1,p}$ with m new unary predicates. Now there is a natural 1-1 correspondence between $\tilde{\mathcal{T}}_{d-1,p}$ and $\mathcal{T}_{d-1,p \cdot 2^m}$, and two $\tilde{\mathcal{T}}_{d-1,p}$ structures are MSO[m]-equivalent iff their corresponding $\mathcal{T}_{d-1,p \cdot 2^m}$ structures are. Then the rank of α_Ψ is $O((\iota(m, \mathcal{T}_{d-1,p \cdot 2^m}, \text{MSO}))^{m+1})$. Let c_0 be a constant such that $\text{rank}(\alpha_\Psi) \leq c_0 \cdot (\iota(m, \mathcal{T}_{d-1,p \cdot 2^m}, \text{MSO}))^{m+1}$ for all $m \geq m_0$. In fact, as shown in Section 2.2, the values of c_0 and m_0 can be taken to be $c_0 = 2$ and $m_0 = 1$.

Returning to α_Φ , we get that $\text{rank}(\alpha_\Phi) \leq c_0 \cdot (\iota(m, \mathcal{T}_{d-1,p \cdot 2^m}, \text{MSO}))^{m+1}$. Now by induction hypothesis, we have that (i) if $d-1 = 0$, then $\iota(m, \mathcal{T}_{d-1,p \cdot 2^m}, \text{MSO}) = p \cdot 2^m$, and (ii) if $d > 1$, then $\iota(m, \mathcal{T}_{d-1,p \cdot 2^m}, \text{MSO}) \leq \zeta_{d-1,p \cdot 2^m}(m, d) = \text{tower}(d, g(d-1) \cdot (m+1) \cdot (2m + \log p)) \leq \text{tower}(d, 2 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$. Then the rank of α_Φ is at most

1. $c_0 \cdot (p \cdot 2^m)^{m+1} \leq 2^{c_0 \cdot (m+1) \cdot (m + \log p)}$ if $d = 1$, and
2. $c_0 \cdot \text{tower}(d, 2 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))^{m+1} \leq \text{tower}(d, c_0 \cdot 4 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$, if $d > 1$.

If $\rho_{d,p} : \mathbb{N} \rightarrow \mathbb{N}$ is the function given by $\rho_{d,p}(m) = \text{tower}(d+1, 4 \cdot c_0 \cdot g(d) \cdot (m+1) \cdot (m + \log p))$, then we see that in either case above, the rank of α_Φ is at most $\rho_{d-1,p}(m)$. We also see that $\tau_{m,\mathcal{T}_{d-1,p}}$ is the vocabulary which contains one unary predicate symbol for every element of $\Delta_{m,\mathcal{T}_{d-1,p}}$, and only those predicates; then $|\tau_{m,\mathcal{T}_{d-1,p}}| = \iota(m, \mathcal{T}_{d-1,p}, \text{MSO})$ which is equal to p if $d-1 = 0$, and at most $\text{tower}(d, g(d-1) \cdot (m+1) \cdot (m + \log p))$ if $d > 1$. Then $\text{rank}(\alpha_\Phi) \cdot |\tau_{m,\mathcal{T}_{d-1,p}}| \leq \rho_{d-1,p}(m) \cdot |\tau_{m,\mathcal{T}_{d-1,p}}| \leq \text{tower}(d, 5 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$ for all $d \geq 1$.

We observe now that $\mathcal{I}_m(\mathcal{F})$ is a structure over a finite monadic vocabulary such that each element of its universe is in the interpretation of exactly one predicate in the vocabulary. Let $\mathbf{s}_{i^*} \in \mathcal{F}$ for $i^* \in I$ be such that the height of \mathbf{s}_{i^*} is equal to $d-1$ (there must be such a tree in \mathcal{F} since height of \mathbf{t} is equal to d). Recall that I is the universe of $\mathcal{I}_m(\mathcal{F})$. Then by Lemma 3.3(1), taking $a = i^*$ and $q = \rho_{d-1,p}(m)$ in the lemma, we get that there exists a substructure \mathcal{B} of $\mathcal{I}_m(\mathcal{F})$ such that (i) \mathcal{B} contains i^* , (ii) $|\mathcal{B}| \leq 1 + (q-1) \cdot |\tau_{m,\mathcal{T}_{d-1,p}}| \leq q \cdot |\tau_{m,\mathcal{T}_{d-1,p}}| \leq \text{tower}(d, 5 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$, and (iii) $\mathcal{B} \equiv_{q,\text{FO}} \mathcal{I}_m(\mathcal{F})$. Then \mathcal{B} can be seen as the MSO[m]-type indicator $\mathcal{I}_m(\mathcal{F}')$ of the family $\mathcal{F}' = (\mathbf{s}_j)_{j \in I'}$ for a subset $I' \subseteq I$, that contains i^* . Then by Theorem 2.1, we have

$$\mathcal{I}_m(\mathcal{F}') \models \alpha_\Phi \quad \text{if, and only if,} \quad \bigcup_{j \in I'} \mathbf{s}_j \models \Phi$$

Since (i) $\mathcal{I}_m(\mathcal{F}') = \mathcal{B} \equiv_{q,\text{FO}} \mathcal{I}_m(\mathcal{F})$, (ii) $\mathcal{I}_m(\mathcal{F}) \models \alpha_\Phi$, and (iii) $\text{rank}(\alpha_\Phi) \leq q$, we have $\mathcal{I}_m(\mathcal{F}') \models \alpha_\Phi$ and therefore $\bigcup_{j \in I'} \mathbf{s}_j \models \Phi$. Then $\bigcup_{j \in I'} \mathbf{s}_j \equiv_{m,\text{MSO}} f$.

We now utilize the induction hypothesis for Part (1). Since $\mathbf{s}_j \in \mathcal{T}_{d-1,p}$ for all $j \in I'$, there exists a leaf-hereditary subtree \mathbf{s}'_j of \mathbf{s}_j such that (i) the heights of \mathbf{s}'_j and \mathbf{s}_j are

the same, (ii) $|s'_j| \leq \zeta_{d-1,p}(m, d-1)$, and (ii) $s'_j \equiv_{m, \text{MSO}} s_j$. Then consider the forest $f' = \cup_{j \in I'} s'_j$. Since disjoint union satisfies the Feferman-Vaught composition property (see Section 2), we get that $f' \equiv_{m, \text{MSO}} \cup_{j \in I'} s_j \equiv_{m, \text{MSO}} f$. Further, we have $1 + |f'| = 1 + \sum_{j \in I'} |s'_j| \leq 1 + |I'| \cdot \max_{j \in I'} |s'_j| \leq 1 + \text{tower}(d, 5 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p)) \cdot \text{tower}(d-1, g(d-1) \cdot (m+1) \cdot (m + \log p)) \leq \text{tower}(d, 6 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$ for all $d \geq 1$. Let t' be the leaf-hereditary subtree of t obtained by removing all child subtrees (that are the unrooted versions of) s_j for $j \notin I'$, and replacing s_j with s'_j (again the replacement being for the unrooted versions of the trees) for each $j \in I'$ in t . Then f' is indeed the forest of rooted trees obtained by deleting the root of t' . Observe that f' contains the tree s'_{i^*} (since $i^* \in I'$) whose height is the same as that of s_{i^*} whose height in turn is equal to $d-1$; then t' has height equal to d which is the height of t . Further, since $f' \equiv_{m, \text{MSO}} f$, we get by Lemma 3.4, that $t' \equiv_{m, \text{MSO}} t$. Finally, $|t'| = |f'| + 1 \leq \text{tower}(d, 6 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$ for all $d \geq 1$.

We now observe that the existence of t' as above for every tree t in $\mathcal{T}_{d,p}$ implies that $\iota(m, \mathcal{T}_{d,p}, \text{MSO})$ is at most the number of structures of $\mathcal{T}_{d,p}$ whose size is at most $\text{tower}(d, 6 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$. Since the number of structures of $\mathcal{T}_{d,p}$ with universe size at most μ for any number μ is at most $\mu \cdot 2^{\mu(\mu + \log p)} \leq 2^{3\mu^2}$ if $\log p \leq \mu$, we get, by taking $\mu = \text{tower}(d, 6 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$, that $\iota(m, \mathcal{T}_{d,p}, \text{MSO}) \leq \text{tower}(d+1, 14 \cdot c_0 \cdot g(d-1) \cdot (m+1) \cdot (m + \log p))$. Then defining $g(d) = 14 \cdot c_0 \cdot g(d-1)$, we see that both parts of the present theorem are true for d . This completes the induction and hence the proof. \square

Remark 3.5. By the same reasoning as in the proof of Theorem 3.2, it follows that if \mathcal{C} in Theorem 2.1 is taken to be $\mathcal{T}_{d,p}$, then for the MSO[m] sentence Φ as considered in Theorem 2.1, so an arbitrary MSO sentence of rank m over the vocabulary of $\mathcal{T}_{d,p}$, the sentence α_Φ is such that (i) the vocabulary $\tau_{m, \mathcal{T}_{d,p}}$ of α_Φ has size at most $\zeta_{d,p}(m, d+1)$, (ii) the rank of α_Φ is at most $\rho_{d,p}(m)$, and (iii) $\text{rank}(\alpha_\Phi) \cdot |\tau_{m, \mathcal{T}_{d,p}}| \leq \zeta_{d+1,p}(m, d+1)$.

The proof of Theorem 3.2 actually shows a stronger relationship between t' and t . We recall the function $\rho_{d,p} : \mathbb{N} \rightarrow \mathbb{N}$ introduced in the proof of Theorem 3.2, defined as $\rho_{d,p}(m) = \text{tower}(d+1, 4 \cdot g(d) \cdot (m+1) \cdot (m + \log p))$ where g is the function given by Theorem 3.2. We now have the following definition.

Definition 3.6. Let $d, p, m \in \mathbb{N}$ be given. For $t_1, t_2 \in \mathcal{T}_{d,p}$ both of the same height, the binary relation $t_1 \leq_m t_2$ is defined inductively on the height $h \in \{0, \dots, d\}$ of the trees as follows:

1. If $h = 0$, then $t_1 \leq_m t_2$ if $t_1 \cong t_2$.
2. Assume \leq_m has been defined when the height of t_1 and t_2 is $h \leq k < d$. Suppose now that t_1 and t_2 have height equal to $k+1$. For $i \in \{1, 2\}$, let f_i be the forest of rooted trees of $\mathcal{T}_{d-1,p}$ obtained by removing the root of t_i and let \mathcal{F}_i be the family of trees constituting f_i . Let $\mathcal{A}_i = \mathcal{I}_m(\mathcal{F}_i)$ be the MSO[m]-type indicator for \mathcal{F}_i .

We now say that $t_1 \leq_m t_2$ holds if: (i) the roots of t_1 and t_2 are the same, and have the same labels; (ii) $\mathcal{A}_1 \equiv_{q, \text{FO}} \mathcal{A}_2$ where $q = \rho_{k,p}(m)$; and (iii) for each tree $t'_1 \in \mathcal{F}_1$, there exists a tree $t'_2 \in \mathcal{F}_2$ such that $t'_1 \leq_m t'_2$.

The tree \mathbf{t}'_2 in the above definition is clearly unique for \mathbf{t}'_1 . We now observe the following properties of the \leq_m relation, the first two of which can be verified by an easy induction, and the last of which can be verified by a very similar reasoning as in the proof of Theorem 3.2.

Lemma 3.7. *For $d, p, m \in \mathbb{N}$, the following are true of the \leq_m relation on $\mathcal{T}_{d,p}$:*

1. \leq_m is reflexive and transitive.
2. Condition (2.i) in the \leq_m definition can be replaced with “ \mathbf{t}_1 is a leaf-hereditary subtree of \mathbf{t}_2 ” to get an equivalent definition.
3. For trees $\mathbf{t}_1, \mathbf{t}_2 \in \mathcal{T}_{d,p}$, if $\mathbf{t}_1 \leq_m \mathbf{t}_2$, then (i) $\mathbf{t}_1 \leq_{m'} \mathbf{t}_2$ for all $m' \leq m$, and (ii) $\mathbf{t}_1 \equiv_{m, \text{MSO}} \mathbf{t}_2$.

We now obtain the following by an entirely analogous set of arguments as for Theorem 3.2.

Corollary 3.8. *For every $d, p, m \in \mathbb{N}$ and every tree $\mathbf{t} \in \mathcal{T}_{d,p}$, there exists a tree $\mathbf{t}' \in \mathcal{T}_{d,p}$ such that (i) $\mathbf{t}' \leq_m \mathbf{t}$, and (ii) $|\mathbf{t}'| \leq \zeta_{d,p}(m, d)$ where $\zeta_{d,p}$ is the function given by Theorem 3.2.*

Before we proceed, we define the following functions $\chi_{d,p}, \xi_{d,p} : \mathbb{N} \rightarrow \mathbb{N}$ for $d \geq 1$. Observe that both of the functions are strictly increasing.

$$\begin{aligned} \chi_{d,p}(m) &= \text{tower}(d+1, g(d+1) \cdot (m+1) \cdot (m + \log p)) \quad (= \zeta_{d+1,p}(m, d+1)) \\ \xi_{d,p}(m) &= \text{tower}(d+1, g(d+1) \cdot (d+1) \cdot (m+1) \cdot (m + \log p)) \end{aligned} \quad (1)$$

An important consequence of Theorem 2.1 and Theorem 3.2 is the following lemma which will be of much use to us in the next section.

Lemma 3.9. *Let $d, p, m \in \mathbb{N}$ be given. Let $\mathcal{F} = (\mathbf{s}_i)_{i \in I}$ be a family of trees of $\mathcal{T}_{d,p}$ where I is an index set of arbitrary cardinality. Let $\mathcal{A} = \mathcal{I}_m(\mathcal{F})$ be the MSO[m] type indicator of \mathcal{F} and let $i^* \in I$ be a given element of \mathcal{A} . Then the following are true for any $q \in \mathbb{N}$:*

1. If $|I| > \chi_{d,p}(m)$, then for every cardinal λ such that $\chi_{d,p}(m) \leq \lambda \leq |I|$, there exists a subset $I' \subseteq I$ such that (i) I' contains i^* , (ii) $|I'| = \lambda$, and (iii) if $\mathcal{F}' = (\mathbf{s}_i)_{i \in I'}$ and $\mathcal{B} = \mathcal{I}_m(\mathcal{F}')$, then $\mathcal{B} \equiv_{q, \text{FO}} \mathcal{A}$ where $q = \rho_{d,p}(m)$.
2. If $|I| > \chi_{d,p}(m)$, then for every cardinal λ such that $\lambda \geq |I|$, there exists a set I' of which I is a subset and trees $\mathbf{s}_i \in \mathcal{T}_{d,p}$ for $i \in I' \setminus I$ such that (i) $|I'| = \lambda$, and (ii) if $\mathcal{F}' = (\mathbf{s}_i)_{i \in I'}$ and $\mathcal{B} = \mathcal{I}_m(\mathcal{F}')$, then $\mathcal{B} \equiv_{q, \text{FO}} \mathcal{A}$ where $q = \rho_{d,p}(m)$. Further, there exists $j \in I$ such that all the trees \mathbf{s}_i for $i \in I' \setminus I$ can be taken to be isomorphic to \mathbf{s}_j .

Proof. The proof is very similar to the proof of Part (1) of Theorem 3.2. Let \mathbf{f} be the forest given by $\mathbf{f} = \bigcup_{i \in I} \mathbf{s}_i$. As in the proof of Theorem 3.2, let \mathcal{S} be the class of rooted forests whose constituent trees are all in $\mathcal{T}_{d,p}$; then $\mathbf{f} \in \mathcal{S}$. Consider the sentence $\Phi := \Theta_\delta$ for $\delta = \delta_{m, \mathcal{S}}^{\text{MSO}}(\mathbf{f})$. By Theorem 2.1, there exists an FO sentence α_Φ of over the vocabulary $\tau_{m, \mathcal{T}_{d,p}}$ such that

$$\mathcal{A} \models \alpha_\Phi \quad \text{if, and only if,} \quad \mathbf{f} \models \Phi$$

By Remark 3.5, we have that $\text{rank}(\alpha_\Phi) \leq \rho_{d,p}(m)$ and that $\text{rank}(\alpha_\Phi) \cdot |\tau_{m, \mathcal{T}_{d,p}}| \leq \chi_{d,p}(m)$. Once again, we observe that \mathcal{A} is a structure over a finite monadic vocabulary such that each element of its universe is in the interpretation of exactly one predicate in the vocabulary. Then by Lemma 3.3, taking $q = \rho_{d,p}(m)$, we get that for every $\lambda \neq |I|$ and $\lambda \geq \chi_{d,p}(m) > (\text{rank}(\alpha_\Phi) - 1) \cdot |\tau_{m, \mathcal{T}_{d,p}}|$, there exists a structure \mathcal{B} such that (i) \mathcal{B} is a substructure of \mathcal{A} containing i^* if $\lambda < |I|$, and \mathcal{A} is a substructure of \mathcal{B} if $\lambda > |I|$, (ii) $|\mathcal{B}| = \lambda$, and (iii) $\mathcal{B} \equiv_{q, \text{FO}} \mathcal{A}$. If I' is the universe of \mathcal{B} , then \mathcal{B} can be seen as the MSO[m]-type indicator of the family $\mathcal{F}' = (\mathfrak{s}_i)_{i \in I'}$ where if $I' \setminus I \neq \emptyset$ (when $\lambda > |I|$), then \mathfrak{s}_i for $i \in I' \setminus I$ is a tree in $\mathcal{T}_{d,p}$. We see that I' always contains i^* , the size of I' is λ , and as shown above $\mathcal{B} \equiv_{q, \text{FO}} \mathcal{A}$.

To see the last statement of the lemma, we see by Lemma 3.3(2) that for any $T \in \tau_{m, \mathcal{T}_{d,p}}$ such that $|T^{\mathcal{A}}| \geq q$, all the elements of $I' \setminus I$ above can be taken such that they are all in the interpretation of T in \mathcal{B} and in the interpretation of no other predicate of $\tau_{m, \mathcal{T}_{d,p}}$ in \mathcal{B} . In other words, one can choose a tree \mathfrak{s}_j for $j \in I$ such that $j \in T^{\mathcal{A}}$ (so $T = \delta_{m, \mathcal{S}}^{\text{MSO}}(\mathfrak{s}_j)$), and for all $i \in I' \setminus I$, it holds that $i \in T^{\mathcal{B}}$ (so $T = \delta_{m, \mathcal{S}}^{\text{MSO}}(\mathfrak{s}_i)$). In other words, $\mathfrak{s}_i \equiv_{m, \text{MSO}} \mathfrak{s}_j$ for all $i \in I' \setminus I$. In particular then, \mathfrak{s}_i can be taken to be isomorphic to \mathfrak{s}_j for all $i \in I' \setminus I$. \square

4. THE EXTENDED LÖWENHEIM-SKOLEM PROPERTY FOR $\text{TM}_{r,p}(d)$

We prove Theorem 1.1 in this section. We prove the theorem in two parts, first for the case when one of η or λ in the MSO-ELS property described in Section 1 is finite, and the other when both of these cardinals are infinite. We recall that a function $f : \mathbb{N}_+ \rightarrow \mathbb{N}$ is a *scale function* if it is strictly increasing. For a cardinal λ , the λ^{th} scale, denoted $\langle \lambda; f \rangle$, is defined as follows. If λ is finite, then $\langle \lambda; f \rangle$ is the interval $[f(\lambda) + 1, f(\lambda + 1)] = \{j \mid f(\lambda) + 1 \leq j \leq f(\lambda + 1)\}$ for $\lambda > 0$, and $[1, f(1)] = \{j \mid 1 \leq j \leq f(1)\}$ for $\lambda = 0$. If λ is infinite, then $\langle \lambda; f \rangle = \{\lambda\}$.

Theorem 4.1. *Let $d, p \in \mathbb{N}$ be given. There exists a $(d+1)$ -fold exponential scale function $\Upsilon_{r,p,d} : \mathbb{N}_+ \rightarrow \mathbb{N}$ such that the following is true. Let $G \in \text{TM}_{r,p}(d)$ be such that $|G| \in \langle \eta; \vartheta_{r,p,d} \rangle$ for some (possibly infinite) $\eta \geq 1$.*

1. (Downward MSO-ELS) *For all cardinals $\lambda \leq \eta$, if λ is finite, then there exists $H \in \text{TM}_{r,p}(d)$ such that (a) $H \subseteq G$, (b) $|H| \in \langle \lambda; \Upsilon_{r,p,d} \rangle$, and (c) $H \equiv_{\lambda, \text{MSO}} G$.*
2. (Upward MSO-ELS) *For all cardinals $\lambda \geq \eta$, if η is finite, then there exists $H \in \text{TM}_{r,p}(d)$ such that (a) $G \subseteq H$, (b) $|H| \in \langle \lambda; \Upsilon_{r,p,d} \rangle$, and (c) $H \equiv_{\eta, \text{MSO}} G$.*

Proposition 4.2. *Let $d, p \in \mathbb{N}$ be given. Let $G \in \text{TM}_{r,p}(d)$ be such that $|G| = \eta$ for an infinite cardinal η . Then for any infinite cardinal λ , there exists $H \in \text{TM}_{r,p}(d)$ such that (i) $H \subseteq G$ if $\lambda \leq \eta$, and $G \subseteq H$ if $\eta \leq \lambda$, (ii) $|H| = \lambda$, and (iii) $H \equiv_{m, \text{MSO}} G$ for all $m \in \mathbb{N}$.*

Proof of Proposition 4.2. It is known [2] that there exists an FO sentence φ_1 that defines $\text{TM}_r(d)$ over all structures. The property of p -labeling – that the unary predicates P_1, \dots, P_p are mutually exclusive and exhaustive (and allowed to be empty) – is easily defined by an FO sentence φ_2 . Then the sentence $\varphi_1 \wedge \varphi_2$ defines $\text{TM}_{r,p}(d)$ over all structures. Now using the fact that MSO collapses to FO over $\text{TM}_{r,p}(d)$ as shown by Theorem 6.1 in Section 6, the statement of the present proposition is just the classical LS theorem which is known to be true over elementary classes of structures. \square

The proof of Theorem 4.1 proceeds by first showing Theorem 4.3 below which is the analogue of Theorem 4.1 for $\mathcal{T}_{d,p}$. This is the core result which can then be easily transferred to $\text{TM}_{r,p}(d)$ via FO interpretations.

Theorem 4.3. *Let $d, p \in \mathbb{N}$ be given. There exists a $(d+1)$ -fold exponential scale function $\vartheta_{d,p} : \mathbb{N}_+ \rightarrow \mathbb{N}$ such that the following is true. Let $\mathfrak{t} \in \mathcal{T}_{d,p}$ be a given tree such that $|\mathfrak{t}| \in \langle \eta; \vartheta_{d,p} \rangle$ for some (possibly infinite) $\eta \geq 1$.*

1. (Downward) *For all cardinals $\lambda \leq \eta$, if λ is finite, then there exists a tree $\mathfrak{t}' \in \mathcal{T}_{d,p}$ such that (i) $|\mathfrak{t}'| \in \langle \lambda; \vartheta_{d,p} \rangle$, and (ii) $\mathfrak{t}' \preceq_\lambda \mathfrak{t}$.*
2. (Upward) *For all cardinals $\lambda \geq \eta$, if η is finite, then there exists a tree $\mathfrak{t}' \in \mathcal{T}_{d,p}$ such that (i) $|\mathfrak{t}'| \in \langle \lambda; \vartheta_{d,p} \rangle$, and (ii) $\mathfrak{t} \preceq_\eta \mathfrak{t}'$.*

Proof of Theorem 4.1. Let $\vartheta_{d,p}$ be the $(d+1)$ -fold exponential function given by Theorem 4.3. Let $q_0 \in \mathbb{N}$ be such that $q_0 = \max\{\max\{\text{rank}(\Xi_{S,p}) \mid S \subseteq [r]^2 \times [p]\}, \text{rank}(\Omega_{r,p,d})\}$ where $\Omega_{r,p,d}$ is as defined in Section 2. Define $\Upsilon_{r,p,d}$ inductively as follows. Firstly, $\Upsilon_{r,p,d}(1) = \vartheta_{d,p}(q_0 + 1)$. Suppose $\Upsilon_{r,p,d}(i)$ has been defined for $i \geq 1$. Let j be the smallest number such that $\vartheta_{d,p}(j) \geq (d+1) \cdot \Upsilon_{r,p,d}(i)$. Then $\Upsilon_{r,p,d}(i+1) = \vartheta_{d,p}(j+1)$. We verify that $\Upsilon_{r,p,d}$ is increasing and $(d+1)$ -fold exponential, and that $\Upsilon_{r,p,d}(i) \geq \vartheta_{d,p}(i)$ for all $i \geq 1$. We now show that $\Upsilon_{r,p,d}$ is as desired. Let $G \in \text{TM}_{r,p}(d)$ be such that $|G| \in \langle \eta; \Upsilon_{r,p,d} \rangle$. Let $(\mathfrak{s}, S) \in \text{Tree}_{r,p}(d)$ be a tree model for G ; then $G = \Xi_{S,p}(\mathfrak{s})$. Since the number of leaves of \mathfrak{s} equals $|G|$, and the height of \mathfrak{s} is d , we see that $|G| \leq |\mathfrak{s}| \leq d \cdot |G| + 1 \leq (d+1) \cdot |G|$. Let η' be such that $|\mathfrak{s}| \in \langle \eta'; \vartheta_{d,p} \rangle$. We now show Part (1) of the theorem. Part (2) can be done similarly.

Let $\lambda \leq \eta$ be finite. Assume $\lambda < \eta$ since if $\lambda = \eta$, we are done by taking $H = G$. Let λ' be the smallest number such that $\vartheta_{d,p}(\lambda') \geq (d+1) \cdot \Upsilon_{r,p,d}(\lambda)$. Then λ' is finite, and $\lambda' - q_0 \geq \lambda$ because $\Upsilon_{r,p,d}(i) \geq \vartheta_{d,p}(i)$, and $\vartheta_{d,p}$ and $\Upsilon_{r,p,d}$ are both increasing. Further $\lambda' < \eta'$ – this is because $|\mathfrak{s}| \geq |G| \geq \Upsilon_{r,p,d}(\eta) \geq \Upsilon_{r,p,d}(\lambda+1) = \vartheta_{d,p}(\lambda'+1)$ the last of these equalities by definition of $\Upsilon_{r,p,d}$. Then by Theorem 4.3 and Lemma 3.7, there exists a leaf-hereditary subtree \mathfrak{s}' of \mathfrak{s} such that (i) $|\mathfrak{s}'| \in \langle \lambda'; \vartheta_{d,p} \rangle$ and (ii) $\mathfrak{s}' \equiv_{\lambda', \text{MSO}} \mathfrak{s}$. Now since $\vartheta_{d,p}(\lambda') \geq \Upsilon_{r,p,d}(\lambda) \geq \Upsilon_{r,p,d}(1) = \vartheta_{d,p}(q_0 + 1)$, we have $\lambda' \geq q_0 + 1$. Then \mathfrak{s}' models $\Omega_{r,p,d}$ so that $(\mathfrak{s}', S) \in \text{Tree}_{r,p}(d)$. Let $H = \Xi_S(\mathfrak{s}')$. Then since \mathfrak{s}' is a leaf-hereditary subtree of \mathfrak{s} , we have $H \subseteq G$. Further since $\lambda' \geq q_0 + \lambda$ and $\mathfrak{s}' \equiv_{\lambda', \text{MSO}} \mathfrak{s}$, we get that $H \equiv_{\lambda, \text{MSO}} G$ (cf. Section 2.1). Finally, we observe that $|H| \leq |\mathfrak{s}'| \leq (d+1) \cdot |H|$. Since $|\mathfrak{s}'| \leq \vartheta_{d,p}(\lambda'+1) = \Upsilon_{r,p,d}(\lambda+1)$ and since $(d+1) \cdot \Upsilon_{r,p,d}(\lambda) \leq \vartheta_{d,p}(\lambda') < |\mathfrak{s}'|$, we get that indeed $|H| \in \langle \lambda; \Upsilon_{r,p,d} \rangle$. \square

We prove Theorem 4.3 in the remainder of this section. We claim that the function $\vartheta_{d,p}$ in the statement of Theorem 4.3 can be taken to be the function $\vartheta_{d,p}(\lambda) = \xi_{d,p}(\lambda)$ for all finite cardinals $\lambda \geq 1$. We prove this below. A key property available to us with this choice is that

$$\vartheta_{d,p}(\lambda) \geq \vartheta_{d,p}(\lambda-1) + \xi_{d-1,p}(\lambda-1) \quad (2)$$

This is because $\xi_{d-1,p}(\lambda-1) + \xi_{d,p}(\lambda-1) \leq 2 \cdot \xi_{d,p}(\lambda-1)$ and the latter now is easily seen to be $\leq \xi_{d,p}(\lambda)$. We now prove the two parts of Theorem 4.3 separately below. We will need the following two lemmas.

Lemma 4.4. *Let $d, p, m \in \mathbb{N}$ be given. Let $t \in \mathcal{T}_{d,p}$ be a given tree and z be a node of t at height h . Let t' be the tree obtained by replacing the (unrooted version of the) subtree t_z of t with (the unrooted version of) a tree $s \in \mathcal{T}_{h,p}$. If $s \leq_m t_z$, then $t' \leq_m t$.*

Proof. We show that the statement is true for the special case when z is a child of $\text{root}(t)$. The general statement follows by repeated applications of this special case to the subtrees rooted at the nodes appearing along the path from the parent of z to the root of t , and the corresponding subtrees rooted at the nodes appearing along the path from the parent of $\text{root}(s)$ to the root of t' (here we are seeing (the unrooted version of) s as a subtree of t').

Consider t ; let f be the forest of rooted trees obtained by deleting the root of t ; one of these rooted trees is t_z . Let f' be the forest obtained by replacing t_z with s , and keeping the other trees of f intact. Let \mathcal{F} and \mathcal{F}' resp. be the families of the trees in the forests f and f' , and let \mathcal{A} and \mathcal{A}' be the MSO[m]-type indicators for \mathcal{F} and \mathcal{F}' . Since $s \leq_m t_z$, it follows that $s \equiv_{m, \text{MSO}} t_z$, so that $\mathcal{A} \cong \mathcal{A}'$ and hence $\mathcal{A}' \equiv_{q, \text{FO}} \mathcal{A}$ for any q . Also for every tree $x \in \mathcal{F}'$, choose $y \in \mathcal{F}$ as: $y = x$ if $x \neq s$ and $y = t_z$ otherwise; then we see that $x \leq_m y$ (since \leq_m is reflexive; see Lemma 3.7). Further the heights of x and y are the same, so that t and t' have the same height. Then all the conditions for $t' \leq_m t$ being true have been met. \square

Lemma 4.5. *Let $d, p, m \in \mathbb{N}$ be given. For each tree $t \in \mathcal{T}_{d,p}$, if $|t| > \xi_{d,p}(m)$, then there is some node z of t at height say h in t , such that (i) the number of children of z in t is $> \chi_{h,p}(m)$, and (ii) for any node $y \in t_z$ such that $y \neq z$, the number of children of y in t (and hence in t_z) is at most $\chi_{h_y,p}(m)$ where h_y is the height of y in t . (Consequently, $|t_y| \leq \xi_{h_y,p}(m)$).*

Proof. Let t be as in the statement of the lemma. Suppose there is no $z \in t$ such that the number of children of z in t is $> \chi_{h,p}(m)$ where h is the height of z in t . Then $|t| \leq 1 + \chi_{d,p}(m) + \chi_{d,p}(m) \cdot \chi_{d-1,p}(m) + \dots + \chi_{d,p}(m) \cdot \chi_{d-1,p}(m) \cdot \dots \cdot \chi_{1,p}(m) \leq (\chi_{d,p}(m))^{d+1} \leq \xi_{d,p}(m)$, contradicting the premise about size of t . Then there exists a node z of smallest height h such that the number of children of z is $> \chi_{h,p}(m)$; so that for every node $y \neq z$ in the subtree t_z of t rooted at z , it is the case that the number of children of y in t (and hence in t_z) is at most $\chi_{h_y,p}(m)$ where h_y denotes the height of y in t . By a similar calculation as above, it follows that $|t_y| \leq \xi_{h_y,p}(m)$ for all $y \neq z$ in t_z . \square

4.1. THE DOWNWARD DIRECTION

We first show the downward direction of Theorem 4.3. Our proof is along similar lines as the proof of [11, Proposition 6.3]. We first need the following result akin to [11, Lemma 6.2].

Lemma 4.6. *Let $d, p, m \in \mathbb{N}$ be given. Then for every tree $t \in \mathcal{T}_{d,p}$ such that $|t| > \xi_{d,p}(m)$, there exists a tree $s \in \mathcal{T}_{d,p}$ such that (i) $|t| - |s| \leq \xi_{d-1,p}(m)$, and (ii) $s \leq_m t$.*

Proof. Let t be as in the statement of the lemma. We first observe by Lemma 4.5 that there is a node z in t at height $h \leq d$ such that (i) the number of children of z in t is $> \chi_{h,p}(m)$, and (ii) $|t_y| \leq \xi_{h-1,p}(m)$ for all children y of z in t . Consider t_z (which has height equal to h); let f be the forest of rooted trees obtained by deleting the root of t_z .

Then $f = \cup_{i \in I} x_i$ where $x_i \in \mathcal{T}_{h-1,p}$ and $|x_i| \leq \xi_{h-1,p}(m)$, and I is of size $> \chi_{h,p}(m)$ (and I could be infinite as well). Let $i^* \in I$ be such that x_{i^*} has height equal to $h-1$ (there must be such an i^* since the height of t_z is h). Then by Lemma 3.9(1), there exists a subset I' of I such that (i) I' contains i^* , (ii) $|I'| = |I| - 1$, and (iii) if \mathcal{A} and \mathcal{A}' are resp. the MSO[m] type indicators for the families of trees in f and $f' = \cup_{i \in I'} x_i$, then $\mathcal{A} \equiv_{q, \text{FO}} \mathcal{A}'$ for $q = \rho_{h,p}(m) > \rho_{h-1,p}(m)$. Let $I \setminus I' = \{j\}$ and let t'_z be the leaf-hereditary subtree of t_z obtained by deleting (the unrooted version of) x_j from t_z ; then f' is the forest of rooted trees obtained by deleting the root of t'_z . We observe now that t'_z and t_z have the same height h (since f' contains x_{i^*}) and further that $t'_z \leq_m t_z$ is indeed true. (The first two conditions of the \leq_m definition are already shown satisfied above, and since $\mathcal{F}' \subseteq \mathcal{F}$, we use the reflexivity of \leq_m from Lemma 3.7 to see that the last condition of the \leq_m definition is satisfied as well). Let s be the leaf-hereditary subtree of t obtained by replacing (the unrooted version of) t_z with (the unrooted version of) t'_z ; then by Lemma 4.4, we get that $s \leq_m t$. We now observe that $|t| - |s| = |t_z| - |t'_z| = |f| - |f'| = |s_j| \leq \xi_{h-1,p}(m) \leq \xi_{d-1,p}(m)$. Then s is indeed as desired. \square

Proof of Theorem 4.3(1). Suppose $t \in \mathcal{T}_{d,p}$ is such that $|t| \in \langle \eta; \vartheta_{d,p} \rangle$. If $\lambda = \eta$, then taking $t' = t$ we are done; so assume $\lambda < \eta$. To prove the theorem, we consider two cases depending on whether η is finite or infinite.

If η is finite, then it suffices to show the existence of the desired tree t' for $\lambda = \eta - 1$. This is because for all $\mu < \eta$, the \leq_μ relation is transitive and because $t' \leq_{\eta-1} t$ implies $t' \leq_\mu t$ (see Lemma 3.7). So let $\lambda = \eta - 1$. Now since $|t| \in \langle \eta; \vartheta_{d,p} \rangle$, we have that $|t| > \vartheta_{d,p}(\eta) = \xi_{d,p}(\eta) > \xi_{d,p}(\lambda)$. Then by Lemma 4.6, we get that there exists $s \in \mathcal{T}_{d,p}$ such that (i) $|t| - |s| \leq \xi_{d-1,p}(\lambda)$, and (ii) $s \leq_\lambda t$. Now we observe that $|\langle \lambda; \vartheta_{d,p} \rangle| = \vartheta_{d,p}(\lambda + 1) - \vartheta_{d,p}(\lambda) \geq \xi_{d-1,p}(\lambda)$ by the inequality (2). Then either $|s| \in \langle \lambda; \vartheta_{d,p} \rangle$ or $|s| \in \langle \eta; \vartheta_{d,p} \rangle$. In the former case, we are done by taking $t' = s$ and in the latter case, we apply Lemma 4.6 recursively to s and again utilize the transitivity of \leq_λ , until we eventually get the desired tree t' .

If η is infinite, then define $\nu = \vartheta_{d,p}(\lambda + 1) + 1$. Observe that $\chi_{d-1,p}(\lambda) < \nu < \eta$. We show below the following statement:

(\dagger) There exists a finite tree $t'' \in \mathcal{T}_{d,p}$ such that (i) $\nu \leq |t''|$, and (ii) $t'' \leq_\lambda t$.

Assuming (\dagger), we see that $|t''| \in \langle \eta'; \vartheta_{d,p} \rangle$ for some finite $\eta' > \lambda$. Then by the previous case above, there exists a tree $t' \in \mathcal{T}_{d,p}$ such that (i) $|t'| \in \langle \lambda; \vartheta_{d,p} \rangle$, and (ii) $t' \leq_\lambda t''$. Then t' is as desired since \leq_λ is transitive. We now show (\dagger) to complete the proof.

Consider t ; it is such that $|t| = \eta$ where η is infinite (since $\langle \eta; \vartheta_{d,p} \rangle = \{\eta\}$ for such η). Then the set $X = \{z \in t \mid z \text{ has } > \nu \text{ many children in } t\}$ is non-empty. (For if not, then every node in t has at most ν children, implying t is finite, since it has height $\leq d$.) We first show the existence of t'' when $\text{root}(t) \in X$. Let f be the forest of rooted trees obtained by deleting the root of t ; then $f = \cup_{i \in I} x_i$ where $x_i \in \mathcal{T}_{d-1,p}$ and $|I| > \nu$. By Corollary 3.8, for each $i \in I$, there exists a tree $x'_i \in \mathcal{T}_{d-1,p}$ such that (i) $|x'_i| \leq \zeta_{d-1,p}(\lambda, d-1)$, and (ii) $x'_i \leq_\lambda x_i$, and hence $x'_i \equiv_{\lambda, \text{MSO}} x_i$ (by Lemma 3.7). Then consider the forest $f' = \cup_{i \in I} x'_i$. We see that the MSO[λ] type indicators \mathcal{A} and \mathcal{A}' resp. of the families of trees in forests f and f' , are isomorphic, and hence FO[q]-equivalent for any q . Let the height of t be $h \leq d$, and let $i^* \in I$ be such that x_{i^*} , and hence x'_{i^*} , has height $h-1$. Since $\nu > \chi_{d-1,p}(\lambda)$, we have by Lemma 3.9(1), that there exists a subset $I'' \subseteq I$ of size exactly ν and containing i^* , such that if $f'' = \cup_{i \in I''} x'_i$, then the MSO[λ] type indicator \mathcal{A}'' of the family of trees in f'' satisfies $\mathcal{A}'' \equiv_{q, \text{FO}} \mathcal{A}' = \mathcal{A}$ for $q = \rho_{d-1,p}(\lambda) \geq \rho_{h-1,p}(\lambda)$. Then if t'' is the leaf-hereditary

subtree of \mathbf{t} obtained by deleting (the unrooted versions of) all child subtrees x_i for $i \notin I''$, and replacing each x_i for $i \in I''$ with x'_i (the replacements being of the unrooted versions of the trees), then we see that $\mathbf{t}'' \leq_\lambda \mathbf{t}$ is true. Further, observe that $\nu \leq |\mathbf{t}''|$, and that $|\mathbf{t}''| \leq 1 + \sum_{i \in I''} |x'_i| \leq 1 + \nu \cdot \zeta_{d-1,p}(\lambda, d-1) < \omega$, so \mathbf{t}'' is finite. Then \mathbf{t}'' is indeed as desired.

We now consider the case when the $\text{root}(\mathbf{t}) \notin X$. Let $X' = \{z \in X \mid z \text{ has no ancestor } z' \text{ in } \mathbf{t} \text{ such that } z' \in X\}$. By the preceding part, we have for each $z \in X'$, that if \mathbf{t}_z is the subtree of \mathbf{t} rooted at z , then there is a finite tree $\mathbf{t}''_z \in \mathcal{T}_{d-1,p}$ such that (i) $\nu \leq |\mathbf{t}''_z|$, and (ii) $\mathbf{t}''_z \leq_\lambda \mathbf{t}_z$. Let \mathbf{t}'' be the leaf-hereditary subtree of \mathbf{t} obtained by replacing (the unrooted version of) the subtree \mathbf{t}_z in \mathbf{t} with (the unrooted version of) \mathbf{t}''_z for all $z \in X'$. It is immediate that $\nu \leq |\mathbf{t}''|$. Further we observe that every node y of \mathbf{t}'' has only finitely many children. This is clear if $y \in \mathbf{t}''_z$ for some $z \in X'$, since \mathbf{t}''_z is itself finite. Else, consider a node y of \mathbf{t}'' that is not an element of \mathbf{t}''_z for any $z \in X'$. If y does not have finitely many children, then $y \in X$ and therefore has an ancestor y' (possibly itself) that is in X' . Then $y \in \mathbf{t}''_{y'}$ (since $y \in \mathbf{t}''$) contradicting the assumption above about y . Thus every node of \mathbf{t}'' has finitely many children; then \mathbf{t}'' is finite since it has height $\leq d$. Further since $\mathbf{t}''_z \leq_\lambda \mathbf{t}_z$ for all $z \in X'$, we get by repeated applications of Lemma 4.4 and the transitivity of \leq_λ from Lemma 3.7, that $\mathbf{t}'' \leq_\lambda \mathbf{t}$, completing the proof of (\dagger) . \square

4.2. THE UPWARD DIRECTION

We now show the upward direction of Theorem 4.3.

Proof of Theorem 4.3(2). If $\lambda = \eta$, then taking $\mathbf{t}' = \mathbf{t}$ we are done. So assume $\lambda > \eta$.

Consider a tree $\mathbf{t} \in \mathcal{T}_{d,p}$ and suppose $|\mathbf{t}| \in \langle \eta; \vartheta_{d,p} \rangle$ for some finite $\eta \geq 1$. Then $|\mathbf{t}| > \vartheta_{d,p}(\eta) = \xi_{d,p}(\eta)$ and hence by Lemma 4.5, there is some node z of \mathbf{t} at height $h \leq d$ such that (i) the number of children of z in \mathbf{t} is $> \chi_{h,p}(\eta)$, and (ii) $|\mathbf{t}_y| \leq \xi_{h-1,p}(\eta)$ for each child y of z .

Let $\lambda > \eta$ be a cardinal. Consider the subtree \mathbf{t}_z of \mathbf{t} rooted at z where z is as above. We show below the following statement:

(\ddagger) There exists a tree $\mathbf{t}'_z \in \mathcal{T}_{h,p}$ of which \mathbf{t}_z is a leaf-hereditary subtree such that (i) if λ is finite, then $|\mathbf{t}'_z| - |\mathbf{t}_z| \leq \xi_{d-1,p}(\eta)$, and if λ is infinite, then $|\mathbf{t}'_z| - |\mathbf{t}_z| = \lambda$, and (ii) $\mathbf{t}_z \leq_\eta \mathbf{t}'_z$.

Then if \mathbf{t}'' is the tree obtained by replacing the subtree \mathbf{t}_z in \mathbf{t} with \mathbf{t}'_z (the replacements being of the unrooted versions of the mentioned trees), then by Lemma 4.4, we have that $\mathbf{t} \leq_\eta \mathbf{t}''$. We now examine $|\mathbf{t}''|$; observe that $|\mathbf{t}''| = |\mathbf{t}'_z| - |\mathbf{t}_z| + |\mathbf{t}|$. We have two cases as below:

1. If λ is finite, then we observe firstly that it suffices to show the theorem for simply $\lambda = \eta + 1$; this is once again due to the transitivity of \leq_η and because \leq_μ implies \leq_η for all $\mu \geq \eta$ (Lemma 3.7). So let $\lambda = \eta + 1$. Then by (\ddagger) , we have $|\mathbf{t}''| - |\mathbf{t}| = |\mathbf{t}'_z| - |\mathbf{t}_z| \leq \xi_{d-1,p}(\eta)$. Now observe that by inequality (2), we have $|\langle \lambda; \vartheta_{d,p} \rangle| = \vartheta_{d,p}(\lambda + 1) - \vartheta_{d,p}(\lambda) \geq \xi_{d-1,p}(\lambda) \geq \xi_{d-1,p}(\eta)$. Then either $\mathbf{t}'' \in \langle \lambda; \vartheta_{d,p} \rangle$ or $\mathbf{t}'' \in \langle \eta; \vartheta_{d,p} \rangle$. In the former case, we are done by taking $\mathbf{t}' = \mathbf{t}''$, and in the latter case, we repeat all of the above arguments with \mathbf{t}'' in place of \mathbf{t} recursively to eventually obtain \mathbf{t}' such that $|\mathbf{t}'| \in \langle \lambda; \vartheta_{d,p} \rangle$. The transitivity of \leq_η again ensures that $\mathbf{t} \leq_\eta \mathbf{t}'$, as desired.
2. If λ is infinite, then by (\ddagger) , $|\mathbf{t}''| = |\mathbf{t}| + \lambda = \lambda$, since \mathbf{t} is finite (as $|\mathbf{t}| \in \langle \eta; \vartheta_{d,p} \rangle$). We are then done by taking $\mathbf{t}' = \mathbf{t}''$.

We now show (‡) to complete the proof.

Let f be the forest of rooted trees obtained by removing the root of t_z . Then $f = \cup_{i \in I} s_i$ where $s_i \in \mathcal{T}_{h-1,p}$ and $|s_i| \leq \xi_{h-1,p}(\eta)$ for all $i \in I$, and $|I| > \chi_{h,p}(\eta)$. By Lemma 3.9(2), for each $\mu > |I|$, there exists $I' \supseteq I$ of size μ and trees $s_i \in \mathcal{T}_{h-1,p}$ for $i \in I' \setminus I$ such that if $f' = \cup_{i \in I'} s_i$, and \mathcal{A} and \mathcal{A}' are resp. the MSO[m] type indicators of f and f' , then $\mathcal{A} \equiv_{q, \text{FO}} \mathcal{A}'$ where $q = \rho_{h,p}(\eta) > \rho_{h-1,p}(\eta)$. Further, for all $i \in I' \setminus I$, the tree s_i is isomorphic to s_j for some $j \in I$. Let now $t'_{z,\mu}$ be the extension of t_z obtained by adding for each $i \in I' \setminus I$, the (unrooted version of) tree s_i as a child subtree of the root of t_z . So that t_z is a leaf-hereditary subtree of $t'_{z,\mu}$, and f' is exactly the forest of rooted trees obtained by deleting the root of $t'_{z,\mu}$. Then t_z and $t'_{z,\mu}$ have the same height h , and further one verifies that $t_z \preceq_\eta t'_{z,\mu}$.

We now choose μ suitably to show that (‡) is true with $t'_z = t'_{z,\mu}$. If λ is finite, then choose $|I'| = \mu = |I| + 1$. Then $|t'_z| - |t_z| = |f'| - |f| = \sum_{i \in I' \setminus I} |s_i| = |s_j| < \xi_{h-1,p}(\eta) \leq \xi_{d-1,p}(\eta)$. If λ is infinite, then choose $|I'| = \mu = \lambda$. Then $|t'_z| - |t_z| = \sum_{i \in I' \setminus I} |s_i| = \lambda \cdot |s_j| = \lambda$. \square

5. AN MSO COMPACTNESS THEOREM

For $r, p \in \mathbb{N}$, recall that every tree model $(s, S) \in \text{Tree}_{r,p}(d)$ can be seen as a tree in $\mathcal{T}_{d,r,p+1}$ via the 1-1 function f that maps the pair (i, j) labeling any leaf node of s with $1 \leq i \leq r$ and $1 \leq j \leq p$, to the number $(i-1) \cdot p + j \in [r \cdot p]$; and every internal node of s is regarded as having the label $r \cdot p + 1$. Let θ be an MSO sentence over the vocabulary of $\mathcal{T}_{d,r,p+1}$ and let $\mathcal{S} \in \mathcal{P}([r]^2 \times [p])$ be a set of signatures, where $\mathcal{P}(X)$ denotes the power set of X . We denote by $\text{Tree}_{r,p}(d; \theta, \mathcal{S})$ the subclass of $\text{Tree}_{r,p}(d)$ consisting of those tree models (s, S) such that $S \in \mathcal{S}$, and s , seen as a tree in $\mathcal{T}_{d,r,p+1}$ as described above, models θ . We correspondingly denote by $\text{TM}_{r,p}(d; \theta, \mathcal{S})$ the subclass of $\text{TM}_{r,p}(d)$ consisting of those graphs that have a tree model in $\text{Tree}_{r,p}(d; \theta, \mathcal{S})$. So $\text{TM}_{r,p}(d) = \text{TM}_{r,p}(d, \text{True}, \mathcal{P}([r]^2 \times [p]))$. We now have the following compactness property for $\text{TM}_{r,p}(d; \theta, \mathcal{S})$.

Theorem 5.1. *Let $d, r, p \in \mathbb{N}$ be given. Let θ be an MSO sentence over the vocabulary of $\mathcal{T}_{d,r,p+1}$ and let $\mathcal{S} \in \mathcal{P}([r]^2 \times [p])$ be a set of signatures. Then for every set $T = \{\varphi_1, \varphi_2, \dots\}$ of MSO sentences over the vocabulary of $\text{TM}_{r,p}(d)$, if every finite subset of T is satisfiable over $\text{TM}_{r,p}(d; \theta, \mathcal{S})$, then T is satisfiable over $\text{TM}_{r,p}(d; \theta, \mathcal{S})$. Further T has a countable model in $\text{TM}_{r,p}(d; \theta, \mathcal{S})$.*

As with earlier results, Theorem 5.1 is established by first showing a corresponding compactness theorem for $\mathcal{T}_{d,p}$, namely Theorem 5.2 below, and then transferring the result via FO interpretations.

Theorem 5.2. *Let $d, p \in \mathbb{N}$ be given. Then for every set $T = \{\varphi_1, \varphi_2, \dots\}$ of MSO sentences over the vocabulary of $\mathcal{T}_{d,p}$, if every finite subset of T is satisfiable over $\mathcal{T}_{d,p}$, then T is satisfiable over $\mathcal{T}_{d,p}$. Further T has a countable model in $\mathcal{T}_{d,p}$.*

Proof of Theorem 5.1. Let $\psi_i = \bigwedge_{j=1}^{j=i} \varphi_i$ for $i \in \mathbb{N}$. By the premise, ψ_i is satisfied in a p -labeled graph $G_i \in \text{TM}_{r,p}(d; \theta, \mathcal{S})$. Let $(s_1, S_1) \in \text{Tree}_{r,p}(d)$ be a tree model for G_i ; then s_i seen as a tree of $\mathcal{T}_{d,r,p+1}$ models θ , and $S_i \in \mathcal{S}$ for all $i \geq 1$. Now since \mathcal{S} is finite, it must be that there exist infinitely many indices $i_1 < i_2 < \dots$ such that $S_{i_1} = S_{i_2} = \dots = S$ (say).

Then consider the FO interpretation $\Xi_{S,p}$ as described in Section 2.1; this is such that $G_{i_j} = \Xi_{S,p}(s_{i_j})$ for $j \geq 1$.

Consider now the theory $T' = \{\Xi_{S,p}(\varphi_i) \mid i \geq 1\} \cup \{\theta, \Omega_{r,p,d}\}$ and the sentences $\psi'_i = \theta \wedge \Omega_{r,p,d} \wedge \bigwedge_{j=1}^{j=i} \Xi_{S,p}(\varphi_i)$ for $i \geq 1$, where $\Omega_{r,p,d}$ and $\Xi_{S,p}(\varphi_i)$ are as defined in Section 2.1. By the fundamental property of interpretations (cf. Section 2.1), we have that $G_{i_j} \models \varphi_l$ iff $s_{i_j} \models \Xi_{S,p}(\varphi_l)$ for all $l, j \geq 1$. Therefore since $G_{i_j} \models \psi'_{i_j}$, we verify that $s_{i_j} \models \psi'_{i_j}$. Then every finite subset of T' is satisfiable in $\mathcal{T}_{d,r,p+1}$. By Theorem 5.2, we have that there is a countable model s^* of T in $\mathcal{T}_{d,r,p+1}$. Since $s^* \models \theta \wedge \Omega_{r,p,d}$ and $S \in \mathcal{S}$, we have that $(s^*, S) \in \text{Tree}_{r,p}(d; \theta, \mathcal{S})$. Let $G^* = \Xi_{S,p}(s^*)$. Then $G^* \in \text{TM}_{r,p}(d; \theta, \mathcal{S})$. Finally, since $s^* \models \Xi_{S,p}(\varphi_i)$ for all $i \geq 1$, we get by the fundamental property of interpretations again, that $G^* \models \varphi_i$ for all $i \geq 1$; then $G^* \models T$. Observe that since s^* is countable, so is G^* . \square

We now devote ourselves to proving Theorem 5.2.

Define a *chain* of trees of $\mathcal{T}_{d,p}$ as a sequence t_1, t_2, \dots of trees of $\mathcal{T}_{d,p}$ such that t_i is a subtree of t_{i+1} for all $i \geq 1$. We denote such a chain \mathcal{C} as $\mathcal{C} := t_1 \subseteq t_2 \subseteq \dots$. We can now define the union t^* of \mathcal{C} , denoted $\bigcup_{i \geq 1} t_i$, in the natural way as: the universe of t^* is the union of the universes of the t_i 's for $i \geq 1$; and for any r -ary predicate R in the vocabulary of $\mathcal{T}_{d,p}$, the interpretation of R in t^* is the union of the interpretations of R in the t_i 's for $i \geq 1$. It is easy to see that t^* belongs to $\mathcal{T}_{d,p}$. We now have the following lemma that plays a crucial role in the proof of Theorem 5.2.

Lemma 5.3. *Let $d, p \in \mathbb{N}$ be given. Let \mathcal{C} be an infinite chain of trees of $\mathcal{T}_{d,p}$ given by $\mathcal{C} := t_1 \subseteq t_2 \subseteq \dots$. Let $t^* = \bigcup_{i \geq 1} t_i$ be the union of \mathcal{C} . Suppose $t_i \leq_{m_i} t_{i+1}$ for all $i \geq 1$, where $0 \leq m_1 \leq m_2 \leq \dots$ and $\sup\{m_i \mid i \geq 1\} = \omega$. Then $t_i \leq_{m_i} t^*$ for all $i \geq 1$.*

Proof of Lemma 5.3. We prove the lemma by induction on d . The base case of $d = 0$ is evident. Assume the statement of the lemma for $d - 1$ where $d \geq 1$. Consider an infinite chain \mathcal{C} of trees of $\mathcal{T}_{d,p}$ given by $\mathcal{C} := t_1 \subseteq t_2 \subseteq \dots$ satisfying the conditions in the statement of the lemma. We first observe that there is no loss of generality in assuming that all the trees of the chain have the same height. For, if the trees have different heights, all of these being $\leq d$ implies that for some $h \leq d$, there is an infinite subchain \mathcal{C}' of \mathcal{C} given by $\mathcal{C}' := t_{j_1} \subseteq t_{j_2} \subseteq \dots$ such that $j_1 < j_2 < \dots$, and all trees in \mathcal{C}' have height h . Then the unions of \mathcal{C}' and \mathcal{C} are the same, namely t^* . Further for $l \geq 1$, since $t_{j_l} \leq_{m_{j_l}} t_{j_{l+1}} \leq_{m_{j_{l+1}}} \dots \leq_{m_{j_{l+1}}} t_{j_{l+1}}$, we infer using Lemma 3.7 that $t_{j_l} \leq_{m_{j_l}} t_{j_{l+1}}$ for all $l \geq 1$. Observe also that $m_{j_l} \leq m_{j_{l+1}}$ for all $l \geq 1$, and $\sup\{m_{j_l} \mid l \geq 1\} = \omega$. Then assuming we have shown the lemma when all trees of \mathcal{C} have equal height, we have $t_{j_l} \leq_{m_{j_l}} t^*$. Then for a given $i \geq 1$, let j_k be such that $i \leq j_k$. Then $t_i \leq_{m_i} t_{j_k} \leq_{m_{j_k}} t^*$; whereby $t_i \leq_{m_i} t^*$ (by Lemma 3.7) showing the lemma for the chain \mathcal{C} . So we henceforth assume that all trees of \mathcal{C} have equal height $h \leq d$.

It is easy to verify that t_i is a leaf hereditary subtree of t^* for all $i \geq 1$. For each tree t_i for $i \geq 1$, consider the forest f_i of rooted trees of $\mathcal{T}_{h-1,p}$ obtained by removing the root of t_i . Likewise let f^* be the forest of rooted trees of $\mathcal{T}_{h-1,p}$ obtained by removing the root of t^* ; then $f^* = \bigcup_{i \geq 1} f_i$. Let \mathcal{F}_i be the family of trees constituting forest f_i and \mathcal{F}^* the family of trees constituting forest f^* , and let $\mathcal{A}_i = \mathcal{I}_{m_i}(\mathcal{F}_i)$ and $\mathcal{A}^* = \mathcal{I}_{m_i}(\mathcal{F}^*)$ be resp. the MSO[m_i] type indicators of \mathcal{F}_i and \mathcal{F}^* .

Fix an $i \geq 1$. For $j \geq i$, since $t_j \leq_{m_j} t_{j+1}$, we have from Lemma 3.7 that $t_j \leq_{m_i} t_{j+1}$ since $m_i \leq m_j$. Then for all $j \geq i$, we have that $\mathcal{A}_j \equiv_{q, \text{FO}} \mathcal{A}_{j+1}$ for $q = \rho_{h-1,p}(m_i)$ (from the \leq_{m_i} definition). We show below that $\mathcal{A}_j \equiv_{q, \text{FO}} \mathcal{A}^*$ for all $j \geq i$. We show this by showing that

for every predicate $T \in \tau_{m_i, \mathcal{T}_{h-1,p}}$, if $|T^{\mathcal{A}_j}| \geq q$, then $|T^{\mathcal{A}^*}| \geq q$ as well, and if $|T^{\mathcal{A}_j}| < q$, then $|T^{\mathcal{A}^*}| = |T^{\mathcal{A}_j}|$. Then $\mathcal{A}_j \equiv_{q, \text{FO}} \mathcal{A}^*$ follows by a simple FO EF game argument.

Let \mathcal{S} be the class of disjoint unions of trees in $\mathcal{T}_{h-1,p}$. Let $|T^{\mathcal{A}_j}| = n$ and let $\mathbf{s}_j^l \in \mathcal{F}_j$ for $1 \leq l \leq n$ be distinct trees such that $\delta_{m_i, \mathcal{S}}(\mathbf{s}_j^l) = T$. Since $\mathbf{t}_i \leq_{m_i} \mathbf{t}_{i+1} \leq_{m_i} \mathbf{t}_{i+2} \leq_{m_i} \dots$, we get that there exist for each l and each $k \geq j$, a unique tree $\mathbf{s}_k^l \in \mathcal{F}_k$ such that $\mathbf{s}_j^l \leq_{m_i} \mathbf{s}_{j+1}^l \leq_{m_i} \mathbf{s}_{j+2}^l \leq_{m_i} \dots$. Then $T = \delta_{m_i, \mathcal{S}}(\mathbf{s}_k^l)$ for all $k \geq j$ and $1 \leq l \leq n$ (since by Lemma 3.7, $\mathbf{s}_j^l \equiv_{m_i, \text{MSO}} \mathbf{s}_{j+1}^l \equiv_{m_i, \text{MSO}} \dots$). Further, since $\mathbf{s}_j^{l_1}$ and $\mathbf{s}_j^{l_2}$ are distinct for $1 \leq l_1 < l_2 \leq n$, it follows that so are $\mathbf{s}_k^{l_1}$ and $\mathbf{s}_k^{l_2}$ for all $k \geq j$. So that $|T^{\mathcal{A}_k}| \geq n$ for all $k \geq j$.

Let's now look at \mathbf{f}^* , \mathcal{F}^* and \mathcal{A}^* . Since $\mathbf{s}_j^l \leq_{m_i} \mathbf{s}_{j+1}^l \leq_{m_i} \mathbf{s}_{j+2}^l \leq_{m_i} \dots$ for $1 \leq l \leq n$, and all trees in the chain belong to $\mathcal{T}_{h-1,p}$ (they could be of different heights), we get by induction hypothesis that $\mathbf{s}_j^l \leq_{m_i} \mathbf{s}_l^* = \bigcup_{k \geq j} \mathbf{s}_k^l$. Then $\delta_{m_i, \mathcal{S}}(\mathbf{s}_l^*) = \delta_{m_i, \mathcal{S}}(\mathbf{s}_j^l) = T$ for all $l \in \{1, \dots, n\}$. Then $|T^{\mathcal{A}^*}| \geq n$ since clearly $\mathbf{s}_{l_1}^*$ and $\mathbf{s}_{l_2}^*$ are distinct for $1 \leq l_1 < l_2 \leq n$. Thus if $|T^{\mathcal{A}_j}| = n \geq q$, then $|T^{\mathcal{A}^*}| \geq q$. Suppose $n < q$ but $|T^{\mathcal{A}^*}| > n$. Then for some $\mathbf{z}^* \in \mathcal{F}^*$, it holds that \mathbf{z}^* is distinct from \mathbf{s}_l^* for all $l \in \{1, \dots, n\}$ and $\delta_{m_i, \mathcal{S}}(\mathbf{z}^*) = T$. Now since $\mathbf{z}^* \in \mathcal{F}^*$, there exists $j^* \geq j$ and trees $\mathbf{z}_k \in \mathcal{F}_k$ for all $k \geq j^*$ such that \mathbf{z}_k is a leaf-hereditary subtree of \mathbf{z}_{k+1} and $\mathbf{z}^* = \bigcup_{k \geq j^*} \mathbf{z}_k$. Since $\mathbf{t}_k \leq_{m_i} \mathbf{t}_{k+1}$, we must have $\mathbf{z}_k \leq_{m_i} \mathbf{z}_{k+1}$ from the definition of \leq_{m_i} , for all $k \geq j^*$. Given that each tree \mathbf{z}_k belongs to $\mathcal{T}_{d-1,p}$, we obtain by our induction hypothesis that $\mathbf{z}_k \leq_{m_i} \mathbf{z}^*$ for each $k \geq j^*$. Then $\mathbf{z}_k \equiv_{m_i, \text{MSO}} \mathbf{z}^*$ (by Lemma 3.7) and hence $\delta_{m_i, \mathcal{S}}(\mathbf{z}_k) = T$ for all $k \geq j^*$ and hence in particular $\delta_{m_i, \mathcal{S}}(\mathbf{z}_{j^*}) = T$. Now since $\mathbf{t}_j \leq_{m_i} \mathbf{t}_{j+1} \leq_{m_i} \dots \leq_{m_i} \mathbf{t}_{j^*}$, we have by the transitivity of \leq_{m_i} that $\mathbf{t}_j \leq_{m_i} \mathbf{t}_{j^*}$ and hence by the definition of \leq_{m_i} , we get that $\mathcal{A}_j \equiv_{q, \text{FO}} \mathcal{A}_{j^*}$. Since $|T^{\mathcal{A}_j}| = n < q$, it follows that $|T^{\mathcal{A}_{j^*}}| = n$ as well. But we now observe that the trees $\mathbf{s}_l^{j^*}$ for $1 \leq l \leq n$ and \mathbf{z}_{j^*} are all distinct, and yet $\delta_{m_i, \mathcal{S}}(\mathbf{s}_l^{j^*}) = T = \delta_{m_i, \mathcal{S}}(\mathbf{z}_{j^*})$ for $1 \leq l \leq n$, which implies $|T^{\mathcal{A}_{j^*}}| > n$ – a contradiction. Then $|T^{\mathcal{A}^*}| = n$.

In summary, we have shown so far conditions (i) and (ii) of the $\mathbf{t}_j \leq_{m_i} \mathbf{t}^*$ definition. To see the last condition, suppose $\mathbf{y}_j \in \mathcal{F}_j$. Then by a similar reasoning as above, there exist trees $\mathbf{y}_k \in \mathcal{F}_k$ for $k \geq j$ such that $\mathbf{y}_j \leq_{m_i} \mathbf{y}_{j+1} \leq_{m_i} \mathbf{y}_{j+2} \leq_{m_i} \dots$. Once again, reasoning as above, the tree $\mathbf{y}^* = \bigcup_{k \geq j} \mathbf{y}_k$ is such that $\mathbf{y}^* \in \mathcal{F}^*$ and $\mathbf{y}_j \leq_{m_i} \mathbf{y}^*$. Thus $\mathbf{t}_j \leq_{m_i} \mathbf{t}^*$ for all $j \geq i$ and hence in particular $\mathbf{t}_i \leq_{m_i} \mathbf{t}^*$. \square

Proof of Theorem 5.2. Let $\psi_i = \bigwedge_{j=1}^{j=i} \varphi_j$ for $i \in \mathbb{N}$. By the premise, we know that ψ_i is satisfiable for all i . Denote by m_i , the rank of ψ_i for all $i \geq 1$; then $1 \leq m_1 \leq m_2 \leq \dots$. If $\sup\{m_i \mid i \geq 1\} < \omega$, then for some $i^* \geq 1$, it holds that $m_j = m_{i^*}$ for all $j \geq i^*$. Then T is a set of MSO formulae of rank at most m_{i^*} , and is hence finite upto equivalence. Then T is already satisfiable in $\mathcal{T}_{d,p}$ since every finite subset of it is satisfiable in $\mathcal{T}_{d,p}$ by assumption. Likewise, if ψ_i has only finitely many models for any $i \geq 1$, then, since ψ_{j+1} implies ψ_j for all $j \geq i$, it holds that for some $j^* > i$ and for all $k \geq j^*$, the sentence ψ_k is equivalent to the (satisfiable) sentence ψ_{j^*} . Then any model of ψ_{j^*} models all of T . We therefore assume below that $\sup\{m_i \mid i \geq 1\} = \omega$, and that ψ_i has infinitely many models for each $i \geq 1$.

Recall the function $\vartheta_{d,p} : \mathbb{N}_+ \rightarrow \mathbb{N}$ that witnesses Theorem 4.3. We observe that for all $i \geq 1$ that $\vartheta_{d,p}(i) \leq \vartheta_{d,p}(i+1)$. Let \mathcal{X}_i for $i \geq 1$ be the class of all finite trees $\mathbf{t} \in \mathcal{T}_{d,p}$ such that (i) $|\mathbf{t}| \in \langle m_i; \vartheta_{d,p} \rangle$, and (ii) $\mathbf{t} \models \psi_i$. Observe that \mathcal{X}_i is a finite class (upto isomorphism). We now claim the following for each $i \geq 1$:

(C1) For every $\mathbf{t} \in \mathcal{X}_{i+1}$, there exists $\mathbf{s} \in \mathcal{X}_i$ such that $\mathbf{s} \leq_{m_i} \mathbf{t}$.

(C2) If \mathcal{X}_i is empty, then ψ_i has only finitely many models in $\mathcal{T}_{d,p}$.

Let us prove the two claims in order. We recall from above that for any $\mathbf{t} \in \mathcal{X}_{i+1}$, we have $|\mathbf{t}| \in \langle m_{i+1}; \vartheta_{d,p} \rangle$. Since $m_{i+1} \geq m_i$, we have by Theorem 4.3(1) that there exists a $\mathbf{s} \in \mathcal{T}_{d,p}$ such that (i) $|\mathbf{s}| \in \langle m_i; \vartheta_{d,p} \rangle$, and (ii) $\mathbf{s} \leq_{m_i} \mathbf{t}$. Now by the latter and Lemma 3.7, we have $\mathbf{s} \equiv_{m_i, \text{MSO}} \mathbf{t}$. Since $\mathbf{t} \in \mathcal{X}_{i+1}$, we have \mathbf{t} models ψ_{i+1} and hence ψ_i (the latter being of rank m_i). Then \mathbf{s} models ψ_i and since $|\mathbf{s}| \in \langle m_i; \vartheta_{d,p} \rangle$, we have $\mathbf{s} \in \mathcal{X}_i$. This shows 1. For 2, suppose that ψ_i has infinitely many models in $\mathcal{T}_{d,p}$. Then by Theorem 4.3(1), there exists a model \mathbf{t} of ψ_i in $\mathcal{T}_{d,p}$ such that $|\mathbf{t}| \in \langle m_i; \vartheta_{d,p} \rangle$; in other words, $\mathbf{t} \in \mathcal{X}_i$ showing 2 in the contrapositive.

Since at the outset itself, we have assumed that all ψ_i 's have infinitely many models, we conclude that $\mathcal{X}_i \neq \emptyset$ for all $i \geq 1$. Further we have already observed above that \mathcal{X}_i is finite for all $i \geq 1$. We now show the following.

(\ddagger): There exists an infinite chain \mathcal{C} given by $\mathcal{C} := \mathbf{t}_1 \subseteq \mathbf{t}_2 \subseteq \dots$ of trees in $\mathcal{T}_{d,p}$ such that for all $i \geq 1$, (i) $\mathbf{t}_i \in \mathcal{X}_i$, and (ii) $\mathbf{t}_i \leq_{m_i} \mathbf{t}_{i+1}$.

We claim that (\ddagger) implies that T is satisfiable over $\mathcal{T}_{d,p}$. Recalling that $\sup\{m_i \mid i \geq 1\} = \omega$ (as assumed at the outset), we have by (\ddagger) and Lemma 5.3, that if $\mathbf{t}^* = \bigcup_{i \geq 1} \mathbf{t}_i$ is the union of the chain \mathcal{C} , then $\mathbf{t}_i \leq_{m_i} \mathbf{t}^*$ for all $i \geq 1$. Then by Lemma 3.7, we have $\mathbf{t}_i \equiv_{m_i, \text{MSO}} \mathbf{t}^*$ for all $i \geq 1$. Since $\mathbf{t}_i \in \mathcal{X}_i$, we have $\mathbf{t}_i \models \psi_i$ and hence $\mathbf{t}^* \models \psi_i$ as well since $\text{rank}(\psi_i) = m_i$. Since this holds for all $i \geq 1$, we have $\mathbf{t}^* \models T$. Further since \mathbf{t}_i is finite for each $i \geq 1$, it follows that \mathbf{t}^* is countable.

We now show (\ddagger) to complete the proof. Let $\text{iso}(\mathcal{X}_i)$ denote the finite set of isomorphism classes of the trees of \mathcal{X}_i for $i \geq 1$.

We construct a labeled tree V with countably many ‘‘levels’’. The root of the tree is at level 0 and for $i \geq 1$, the i^{th} level contains $|\text{iso}(\mathcal{X}_i)|$ many nodes. Each node at level i is labeled with exactly one element of $\text{iso}(\mathcal{X}_i)$ and no two nodes at level i get the same label. (It is possible though that the same label appears in different levels.) This defines the vertices and their labeling. The edge relation is now defined as follows. Firstly the relation relates only pairs of vertices in adjacent levels. All nodes of level 1 are adjacent to the root. For a node u at level $i+1$ for $i \geq 1$, let the isomorphism class of $\mathbf{t} \in \mathcal{X}_{i+1}$ be the label of u . By Claim 1, there exists some node v at level i and some tree $\mathbf{s} \in \mathcal{X}_i$ such that the isomorphism class of \mathbf{s} is the label of v , and $\mathbf{s} \leq_{m_i} \mathbf{t}$. We choose exactly one such node v and put an edge between u and v . The edges as specified above are the only edges in V ; it follows then that V is a tree.

We now observe that V is infinite (since \mathcal{X}_i is non-empty for all i) and finitely branching (since $\text{iso}(\mathcal{X}_i)$ is finite for all i). Then by König’s lemma, there exists an infinite path u_0, u_1, u_2, \dots in V starting at the root u_0 . Let the isomorphism class of \mathbf{t}_i be the label of u_i for $i \geq 1$. Then from the properties of the edge relation, we have for all $i \geq 1$, that (i) $\mathbf{t}_i \in \mathcal{X}_i$, and (ii) $\mathbf{t}_i \leq_{m_i} \mathbf{t}_{i+1}$. Indeed then $\mathcal{C} := \mathbf{t}_1 \subseteq \mathbf{t}_2 \subseteq \dots$ is the desired chain. \square

6. MSO = FO OVER BOUNDED SHRUB-DEPTH CLASSES

We show the following in this section. This result for $p = 1$ is already known from [2], and for $p = 1$ and finite graphs, is known earlier still from [5].

Theorem 6.1. *Let $r, p, d \in \mathbb{N}$ be given. Then over $\text{TM}_{r,p}(d)$, every MSO sentence is equivalent to an FO sentence. Consequently, the same holds over any subclass of $\text{TM}_{r,p}(d)$.*

We need the following notion for the proof. Given any MSO formula $\theta(\bar{x})$, and an FO formula $\gamma(y)$, we can construct a *relativized* MSO formula $\theta|_\gamma(\bar{x})$ that relativizes all quantifiers of θ to range over only elements y satisfying $\gamma(y)$. Specifically, to obtain $\theta|_\gamma(\bar{x})$, we replace every subformula of the form $\exists v\psi(v, \bar{u})$ in $\theta(\bar{x})$ with $\exists v(\gamma(v) \wedge \psi|_\gamma(v, \bar{u}))$ and every subformula of the form $\forall v\psi(v, \bar{u})$ with $\forall v(\gamma(v) \rightarrow \psi|_\gamma(v, \bar{u}))$. The key property of relativization is that for any given structure \mathcal{A} , if \mathcal{A}_γ is the substructure of \mathcal{A} induced by the set $A' = \{b \in \mathcal{A} \mid \mathcal{A} \models \gamma(b)\}$, then for an MSO sentence θ , we have that $\mathcal{A} \models \theta|_\gamma$ iff $\mathcal{A}|_\gamma \models \theta$. In addition to the mentioned notion, we will need the following result from classical model theory [9].

Theorem 6.2. *Let \mathcal{A} and \mathcal{B} be two countable structures over the same vocabulary. If $\mathcal{A} \equiv_{n, \text{FO}} \mathcal{B}$ for all $n \geq 1$, then $\mathcal{A} \cong \mathcal{B}$.*

Proof of Theorem 6.1. Our proof is an adaptation of the proof of Lindström's theorem (mentioned at the outset in Section 1) as presented in [12]. We prove Theorem 6.1 by contradiction. Suppose there is an MSO sentence α over the vocabulary τ of $\text{TM}_{r,p}(d)$ that is not equivalent over $\text{TM}_{r,p}(d)$ to any FO sentence. Then for every $n \in \mathbb{N}$, there exist graphs $G_n, H_n \in \text{TM}_{r,p}(d)$ such that (i) $G_n \equiv_{n, \text{FO}} H_n$, and (ii) $G_n \models \alpha$ but $H_n \models \neg\alpha$. Let H'_n be the graph obtained from H_n by changing every vertex label $j \in [p]$ to $j + p \in [2p]$. This graph is in $\text{TM}_{r,2p}(d)$; take any tree model if $(\mathbf{t}_1, S_1) \in \text{Tree}_{r,p}(d)$ is a tree model for H_n , then changing every leaf label of the form (i, j) in \mathbf{t}_1 to $(i, j + p)$ gives a tree model $(\mathbf{t}'_1, S_1) \in \text{Tree}_{r,2p}(d)$ for H'_n . Let $(\mathbf{t}_2, S_2) \in \text{Tree}_{r,p}(d)$ be a tree model for G_n ; then $(\mathbf{t}_2, S_2) \in \text{Tree}_{r,2p}(d)$ as well. Consider now the tree model $(\mathbf{s}, S) \in \text{Tree}_{r,2p}(d+1)$ where \mathbf{s} is obtained by making (the unrooted versions of) \mathbf{t}_2 and \mathbf{t}'_1 as the (only) child subtrees of a new root node, and $S = S_1 \cup S_2$. It is easily verified that (\mathbf{s}, S) is a tree model for the disjoint union $\mathcal{M}_n = G_n \cup H'_n$ of G_n and H'_n . Observe that $S \subseteq [r]^2 \times [d]$.

Consider now the FO sentence θ over the vocabulary of $\mathcal{T}_{d+1, r \cdot 2p+1}$ that asserts over the subclass $\text{Tree}_{r,2p}(d+1)$ of $\mathcal{T}_{d+1, r \cdot 2p+1}$ that the root has exactly 2 children x_1 and x_2 , and any leaf node whose label corresponds (via the bijective function $f : [r] \times [2p] \rightarrow [r \cdot 2p]$ seen in Section 2.1) to (i) the pair (i, j) where $i \in [r]$ and $j \in [p]$ has a path of length exactly d to x_1 , and (ii) the pair (i, j) where $i \in [r]$ and $j \in \{p+1, \dots, 2p\}$ has a path of length exactly d to x_2 . Let $\mathcal{S} = \mathcal{P}([r]^2 \times [d])$. Then any graph in $\text{TM}_{r,p}(d+1; \theta, \mathcal{S})$ is such that it is a disjoint union of two graphs say A and B where $A \in \text{TM}_{r,p}(d)$ and B is "affinely" in $\text{TM}_{r,p}(d)$ in that, the labels of its vertices belong to $\{p+1, \dots, 2p\}$ and the graph B' obtained by changing every label j of a vertex in B to $j - p$, satisfies $B' \in \text{TM}_{r,p}(d)$. We see that for each $n \geq 1$, the $2p$ -labeled graph \mathcal{M}_n indeed belongs to $\text{TM}_{r,2p}(d+1; \theta, \mathcal{S})$.

We now code the properties of G_n and H_n mentioned above into an MSO sentence β_n over the vocabulary $\tau = \{E, P_1, \dots, P_{2p}\}$ of $\text{TM}_{r,2p}(d+1)$. To be able to do this, define for an MSO sentence ψ over the vocabulary $\{E, P_1, \dots, P_p\}$ of $\text{TM}_{r,p}(d)$, the "affine" MSO sentence $\psi[i \mapsto i + p]$ which is obtained from ψ by replacing every atomic predicate of the form $P_i(z)$ with $P_{i+p}(z)$ for $1 \leq i \leq p$. This sentence is over the vocabulary $\{E, P_{p+1}, \dots, P_{2p}\} \subseteq \tau$. We are now ready to define β_n . Below Δ is a notational short hand for $\Delta_{n, \text{TM}_{r,p}(d)}^{\text{FO}}$. Recall also from Section 2.1, the sentence Θ_δ that defines an equivalence

class $\delta \in \Delta$. Recall also the sentence α from the outset.

$$\begin{aligned} \gamma_1(x) &:= \bigvee_{i=1}^{i=p} P_i(x) \quad ; \quad \gamma_2(x) := \bigvee_{i=p+1}^{i=2p} P_i(x) \\ \beta_n &:= \exists x \gamma_1(x) \wedge \exists x \gamma_2(x) \wedge \\ &\quad \alpha|_{\gamma_1} \wedge (\neg \alpha[i \mapsto i+p])|_{\gamma_2} \wedge \bigwedge_{\delta \in \Delta} \Theta_\delta|_{\gamma_1} \leftrightarrow \Theta_\delta[i \mapsto i+p]|_{\gamma_2} \end{aligned}$$

So β_n asserts that over $\text{TM}_{r,2p}(d+1; \theta, \mathcal{S})$, a $2p$ -labeled graph, which is a disjoint union of two non-empty graphs $A \in \text{TM}_{r,p}(d)$ and B that is affinely in $\text{TM}_{r,p}(D)$ as seen above, is such that (i) A models α but B (seen as a p -labeled graph) does not, and (ii) $A \equiv_{n, \text{FO}} B$ (where B is seen as a p -labeled graph). Clearly, $\mathcal{M}_n \models \beta_n$. Further observe that β_n implies $\beta_{n'}$ for $n \geq n'$.

Let $T = \{\beta_n \mid n \geq 1\}$. Then every finite subset of T is satisfiable $\text{TM}_{r,2p}(d+1; \theta, \mathcal{S})$ since if k is the largest index such that β_k appears in the subset, then the subset is equivalent to β_k and is hence satisfied in $\mathcal{M}_k \in \text{TM}_{r,2p}(d+1; \theta, \mathcal{S})$. By Theorem 5.1, we get that T is satisfied in a countable model $\mathcal{M}^* \in \text{TM}_{r,2p}(d+1; \theta, \mathcal{S})$. Then by our observations above, \mathcal{M}^* is a disjoint union of two graphs G^* and H^{**} such that $G^* \in \text{TM}_{r,p}(d)$ and H^{**} is affinely in $\text{TM}_{r,p}(d)$, that is the graph H^* obtained by changing every label $j \in \{p+1, \dots, 2p\}$ of a vertex in H^{**} , to $j-p$ is such that $H^* \in \text{TM}_{r,p}(d)$. Since $\mathcal{M}^* \models T$, we get that (i) $G^* \models \alpha$ and $H^{**} \models \neg \alpha[i \mapsto i+p]$, so $H^* \models \neg \alpha$, and (ii) $G^* \equiv_{n, \text{FO}} H^*$ for all $n \geq 1$. But now since \mathcal{M}^* is countable, so are G^* and H^* , and therefore by Theorem 6.2, we have $G^* \cong H^*$ – this is a contradiction since G^* models α but H^* does not. \square

7. CONCLUSION

In this paper, we looked at classes of arbitrary graphs of bounded shrub-depth, specifically the class $\text{TM}_{r,p}(d)$ of p -labeled arbitrary graphs whose underlying unlabeled graphs have tree models of height d and r labels. We showed that this class satisfies an extension of the classical Löwenheim-Skolem (LS) property into the finite and for MSO. This extension being a generalization of the small model property, shows the pseudo-finiteness of the graphs of $\text{TM}_{r,p}(d)$. In addition, we obtain as consequences entirely new proofs of various known results concerning bounded shrub-depth classes (of finite graphs) and $\text{TM}_{r,p}(d)$. These include the small model property for MSO with elementary bounds, the classical compactness theorem from model theory over $\text{TM}_{r,p}(d)$ (known because $\text{TM}_{r,p}(d)$ is an elementary class), and the equivalence of MSO and FO over $\text{TM}_{r,p}(d)$ and hence over bounded shrub-depth classes. We prove the last of these by adapting the proof of the classical Lindström's theorem characterizing FO over arbitrary structures.

An interesting consequence of our results is that they allow for transferring results back and forth between finite bounded shrub-depth graphs and their infinite counterparts. For instance, one can obtain a different proof of the MSO-FO equivalence over $\text{TM}_{r,1}(d)$ by simply lifting this equivalence known over bounded shrub-depth classes of finite graphs, to $\text{TM}_{r,1}(d)$. Similarly, one can show that $\text{TM}_{r,1}(d)$ has only finitely many forbidden finite induced subgraphs, where these subgraphs are the same as those constituting a forbidden induced subgraph characterization (in the finite) of the subclass of finite graphs of $\text{TM}_{r,1}(d)$ [3]. (Indeed, the mentioned subgraphs also characterize all of $\text{TM}_{1,p}(d)$ over arbitrary structures, since this class is known to be defined by a finite set of forbidden

finite induced subgraphs [2].) Finally, an example of an infinite to finite transfer is classical preservation theorems, such as the Łoś-Tarski theorem, that is true of $\text{TM}_{r,p}(d)$ since it is an elementary class, and hence true of the class of finite graphs of $\text{TM}_{r,p}(d)$; this is because any FO sentence that defines a hereditary subclass of the latter class also defines a hereditary subclass of the former owing to the fact that the graphs of $\text{TM}_{r,p}(d)$ are pseudo-finite.

For future work, we believe that similar results as in the paper can be obtained for guarded second order logic (GSO) and infinite graphs of bounded tree-depth, since the Feferman-Vaught composition theorem of [4] that we use as our main tool in the paper, has also been shown for GSO. In particular therefore, it might be possible to get an alternative Lindström-style proof as above, of the equivalence of GSO and FO over finite graphs of bounded tree-depth [4]. We would also like to investigate extensions of our results to the extension CMSO of MSO with modulo counting quantifiers. Finally, we would like to explore what further results from classical model theory can be transferred to the finite for bounded shrub-depth classes and conversely, what finitary results can be transported to the infinite.

Acknowledgements: I would like to thank Pascal Schweitzer for discussions pertaining to the \mathcal{L} -ELS property. I also thank Anand Pillay for pointing me to [12] for the proof of Lindström’s theorem.

REFERENCES

- [1] Chen Chung Chang and H Jerome Keisler. *Model theory*. Elsevier, 1990.
- [2] Yijia Chen and Jörg Flum. FO-Definability of Shrub-Depth. In Maribel Fernández and Anca Muscholl, editors, *28th EACSL Annual Conference on Computer Science Logic (CSL 2020)*, volume 152 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [3] Patrice Ossona de Mendez, Jan Obdržálek, Jaroslav Nešetřil, Petr Hliněný, and Robert Ganian. Shrub-depth: Capturing height of dense graphs. *Logical Methods in Computer Science*, 15, 2019.
- [4] Michael Elberfeld, Martin Grohe, and Till Tantau. Where first-order and monadic second-order logic coincide. In *LICS 2012, Croatia, June 25-28, 2012*, pages 265–274, 2012.
- [5] Jakub Gajarsky and Petr Hliněný. Kernelizing MSO properties of trees of fixed height, and some consequences. *Log. Meth. Comp. Sci.*, 11(19):1–26, 2015.
- [6] Robert Ganian, Petr Hliněný, Jaroslav Nešetřil, Jan Obdržálek, Patrice Ossona de Mendez, and Reshma Ramadurai. When trees grow low: shrubs and fast mso 1. In *International Symposium on Mathematical Foundations of Computer Science*, pages 419–430. Springer, 2012.
- [7] O joungh Kwon, R. McCarty, Sang il Oum, and P. Wollan. Obstructions for bounded shrub-depth and rank-depth. *ArXiv*, abs/1911.00230, 2019.

- [8] Leonid Libkin. *Elements of finite model theory*. Springer Science & Business Media, 2013.
- [9] Per Lindström. On extensions of elementary logic. *Theoria*, 35(1):1–11, 1969.
- [10] Abhisekh Sankaran. A finitary analogue of the downward Löwenheim-Skolem property. In Valentin Goranko and Mads Dam, editors, *26th EACSL Annual Conference on Computer Science Logic (CSL 2017)*, volume 82 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 37:1–37:21, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [11] Abhisekh Sankaran. A finitary analogue of the downward Löwenheim-Skolem property. *CoRR*, abs/1705.04493, 2017.
- [12] Jouko Väänänen. Lindström’s theorem. *Universal Logic: An Anthology*, pages 231–236, 2010.