Branchwise-real trees and bisimulations of potentialist systems

Sam Adam-Day

Lady Margaret Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2023

Acknowledgements

I owe a large debt of gratitude to my principal supervisor Joel David Hamkins. Joel supported me throughout my DPhil journey, providing me with invaluable feedback, stimulation, encouragement and professional development opportunities. I am grateful for his flexibility and openness during the latter part of my doctorate, as I transitioned away from set theory. I would also like to thank my secondary supervisor Udi Hrushovski for his guidance and support at various stages of my DPhil.

I wish to thank the anonymous reviewers whose comments greatly improved the research papers incorporated into this thesis. I am also grateful to the EPSRC for its financial support while pursuing my DPhil.

Finally, I thank my friends and family for everything they have done for me throughout the years.

Abstract

This thesis concerns two topics. The first treats $\mathbb{R}$-trees, which are a certain kind of metric space tree in which every point can be branching. Favre and Jonsson posed the following problem in 2004: can the class of partial orders underlying $\mathbb{R}$-trees be characterised by the fact that every branch is order-isomorphic to a real interval? I first answer this question in the negative, then go on to establish a connection between these trees and traditional set-theoretic trees. This connection is then put to work, answering refinements of Favre and Jonsson's question, yielding several independence results. I next move on to consider the existence of examples of these partial orders without non-trivial automorphisms. I provide constructions of these subject to increasingly strong uniformity conditions. While these constructions all take place in ZFC, they have a strong forcing flavour.

The second topic deals with bisimulations of potentialist systems, which are first-order Kripke models based on embeddings. Given a first-order theory T we can impose a potentialist structure on the class of models of T by taking either all embeddings or all substructure inclusions between models. I show that these two systems are always bisimilar. Next, by connecting with a generalisation of the Ehrenfeucht-Fraïsé game, I show the equivalence of the existence of a bisimulation with elementary equivalence with respect to an infinitary language. Finally, I turn to the question of when a class-sized potentialist system is bisimilar to a set-sized one, providing two different sufficient conditions.

Contents

List of Figures

Introduction

This thesis consists of the following three chapters, which I have aimed to publish separately as research papers.

- Chapter 1: "On the continuous gradability of the cut-point orders of $\mathbb{R}$-trees", published in "Topology and its Applications" [Ada22b].
- Chapter 2: "Uniform, rigid branchwise-real trees", to appear in the "Israel Journal of Mathematics" [Ada23].
- Chapter 3: "Bisimulations of potentialist systems", submitted to the "Journal of Symbolic Logic" [Ada22a].

The first two chapters stem from the study of $\mathbb{R}$-*trees*: metric spaces which have a tree structure. Fixing any point p in an $\mathbb{R}$-tree, one can form a partial order tree by considering all geodesic segments beginning at p, ordered by inclusion. These partial orders are called *cut-point orders* and were studied by Favre and Jonsson [FJ04], who gave an axiomatisation (Theorem 1.6). What first piqued my interest in this topic was an open question posed by Favre and Jonsson in this work: whether one component of the axiomatisation — that there exists a continuous monotonic map from the partial order to $\mathbb{R}$ — is necessary. In the first half of Chapter 1 I show that this condition is indeed necessary, by constructing a counterexample using a set-theoretic structure investigated by Baumgartner, Laver and Gavin in 1970. In the second part, I develop the connection with set theory further by providing a necessary and sufficient condition for the existence of a continuous monotonic map into $\mathbb{R}$ in terms of the well-stratified subtrees of the partial order. Finally, in the third section I use this connection to answer certain refinements of the motivating question, some of which turn out to be independent of ZFC.

Chapter 2 investigates these partial orders — which I term *branchwise-real trees* — from a different angle. Motivated by the study of automorphisms of well-stratified trees, as well as the use of $\mathbb{R}$-trees as underlying spaces of group actions in geometric group theory, I look for examples of rigid branchwise-real trees subject to increasingly strong uniformity conditions. I show that there is a rigid branchwise-real tree in which every branching point has the same degree, one in which every point is branching and of the same degree, and finally one in which every point is branching of the same degree and which admits no monotonic map into $\mathbb{R}$. Trees are grown iteratively in stages, and a key technique is the construction of a family of colourings of $(0, \infty)$ which is 'sufficiently generic', using these colourings to determine how to proceed with the construction. While all constructions take place in ZFC, there is a clear connection with forcing, which I elaborate on after each result.

The third chapter deals with a different subject: *potentialist systems*. A potentialist system is a first-order Kripke model based on embeddings, which allows one to study a mathematical structure within the context of a system of related structures. I define the notion of bisimulation for these systems, and provide a number of examples. Given a first-order theory T, the system $\mathrm{Mod}(T)$ consists of all models of T. We can then take either all embeddings, or all substructure inclusions, between these models. While this yields two different potentialist systems, I show that these two ways of defining $\mathrm{Mod}(T)$ are bisimilar. Next, I relate the notion of bisimulation to a generalisation of the Ehrenfeucht-Fraïsé game, and use this to show the equivalence of the existence of a bisimulation with elementary equivalence with respect to an infinitary language. Finally, while potentialist systems can be class-sized in general, they sometimes turn out to be bisimilar with set-sized systems. I consider the question of when this occurs, and give two different sufficient conditions.

Chapter 1

On the continuous gradability of the cut-point orders of $\mathbb{R}$-trees

1.1 Introduction

An $\mathbb{R}$-tree is to the real numbers what a graph-theoretic tree is to the integers. Formally, let $\langle X, d\rangle$ be a metric space. An *arc* between $x, y \in X$ is the image of a topological embedding $r\colon [a, b] \to X$ of a real interval such that $r(a) = x$ and $r(b) = y$ (allowing for the possibility that $a = b$). The arc is a *geodesic segment* if r can be taken to be an isometry. The metric space $\langle X, d\rangle$ is an $\mathbb{R}$-*tree* if between any two points $x, y \in X$ there is a unique arc, denoted $[x, y]$, which is also a geodesic segment.

Note that any tree in the graph-theoretic sense can be viewed as an $\mathbb{R}$-tree via its so-called 'geometric realisation'. Indeed, let V be a (possibly infinite) set of vertices and let E be a symmetric binary relation on V, such that $G = \langle V, E\rangle$ is a connected graph without cycles. Then G can be realised as an $\mathbb{R}$-tree by taking V as a discrete set of points and adding a copy of the unit interval $(0, 1)$ between $u, v \in V$ whenever $\langle u, v\rangle \in E$.

The class of $\mathbb{R}$-trees however is much more general than this. Consider the following example of an $\mathbb{R}$-tree which does not arise in this fashion. Let X be the space obtained by taking the real plane $\mathbb{R}^2$, and designating each point on the x-axis as a 'train station' and each vertical line $\{x\} \times \mathbb{R}$, as well as the x-axis, as a 'train track'. We define a new metric on X: to travel between two points in the real plane, one must travel along the tracks, potentially passing through train stations. In the resulting $\mathbb{R}$-tree, the removal of any point on the x-axis leaves exactly 4 connected components. See Figure 1.1(a) for some example geodesic segments in this tree.

Now, $\mathbb{R}$-trees play an important role in geometric group theory, and are interesting objects in their own right. (See [Bes01; MNO92] and the references contained in

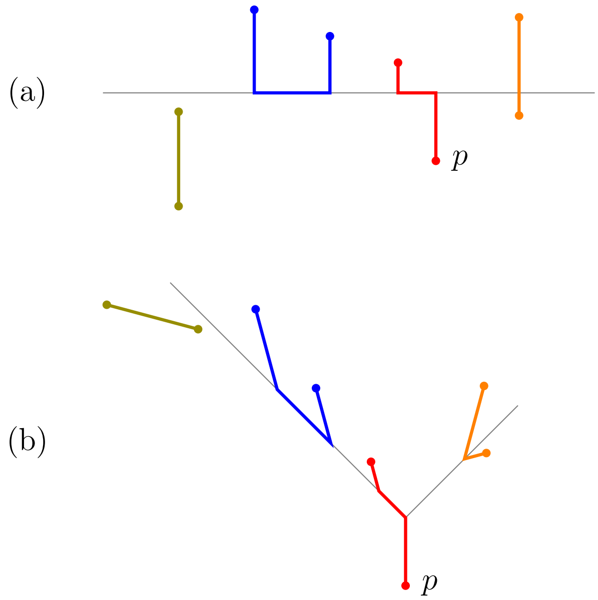

Figure 1.1: (a) Some example geodesic segments in the train track $\mathbb{R}$-tree. (b) The cut-point order on this tree with root p, showing the images of the example geodesic segments.

[Fab15].) In this chapter, we are interested in the underlying order structure of $\mathbb{R}$-trees. Given an $\mathbb{R}$-tree X together with a designated point $p \in X$, the *cut-point order* on X with root p is defined by, for $x, y \in X$:

$$x \leqslant y \quad \Leftrightarrow \quad [p, x] \subseteq [p, y]$$

See Figure 1.1(b) for a picture of a cut-point order on the train track $\mathbb{R}$-tree example.

In §3.1 of [FJ04], Charles Favre and Mattias Jonsson investigate this order structure, one of their aims being an order-theoretic characterisation of structures arising in this way. I replicate their definitions here, with a few changes of terminology.

Definition 1.1. A *tree order* is a partial order X such that the following conditions hold.

(TO1) For every $x \in X$ the set $\downarrow(x) = \{y \in X \mid y \leqslant x\}$ is a linear order.

(TO2) X has a minimum element, its *root*.

A *branch* in X is a maximal linearly ordered subset.

Partial orders satisfying (TO1) are sometimes called 'psuedotrees' (see for instance [Nik89]).

Definition 1.2. A *branchwise-real tree* is a tree order X subject to the following extra conditions.

(BR1) Every branch is order-isomorphic to a real interval.

(BR2) X is a meet-semilattice; that is, any two points $x, y \in X$ have a greatest lower bound $x \wedge y$, their *meet*.

Favre and Jonsson call such objects 'non-metric trees'.

Remark 1.3. In fact, Favre and Jonsson's definition is equivalent to (TO1) + (TO2) + (BR1). They erroneously claim that (BR2) follows from the rest using the completeness of the real line [FJ04, p. 45]. This is incorrect in light of the following counterexample. Let X be the interval $[0, 1)$, together with two incomparable copies of the element 1 sitting on top, as in the following diagram.

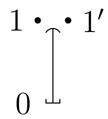

Then 1 and $1'$ have no common meet, and so X is not a branchwise-real tree.

Definition 1.4. Let P and Q be partial orders. A *Q-grading* of P is a strictly monotonic map $f\colon P \to Q$. That is, whenever $x < y$ in P we have $f(x) < f(y)$.

Definition 1.5. Let X be a branchwise-real tree. An $\mathbb{R}$-grading $\ell\colon X \to \mathbb{R}$ is *continuous* if for any $x < y$ in X, letting $[x, y] := \{z \in X \mid x \leqslant z \leqslant y\}$, the restriction:

$$\ell\restriction[x, y]\colon [x, y] \to [\ell(x), \ell(y)]$$

is an order-isomorphism. I will usually drop the '$\mathbb{R}$' and call such functions *continuous gradings*. Say that X is *continuously gradable* if it admits a continuous grading.

In [FJ04] these functions are called 'parametrizations'.

We can now state the order-theoretic characterisation of $\mathbb{R}$-tree cut-point orders which Favre and Jonsson obtained.

Theorem 1.6. *The class of $\mathbb{R}$-tree cut-point orders is exactly the class of continuously gradable branchwise-real trees.*

Proof. See [FJ04, p. 50]. Given an $\mathbb{R}$-tree X and $p \in X$, it is straightforward to verify that the cut-point order on X with root p is a branchwise-real tree. Moreover, we can use the metric d on X to define a continuous grading by:

$$\ell(x) := d(p, x)$$

Conversely, given a branchwise-real tree X and a continuous grading $\ell\colon X \to \mathbb{R}$, we can use ℓ to define a metric on X via the 'railroad track equation' (see also [MNO92]):

$$d(x, y) := \ell(x) + \ell(y) - 2\ell(x \wedge y)$$

See Figure 1.2 for an illustration of railroad track equation. □

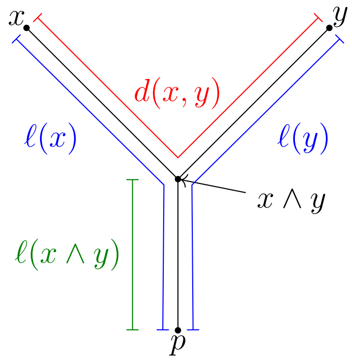

Figure 1.2: The railroad track equation

The construction of the metric on the branchwise-real tree X relies on the continuous grading. It is natural however to imagine that one could always find a continuous grading on a branchwise-real tree using its real-line-like structure. But Favre and Jonsson state: 'We do not know if there exists a non-parameterizable nonmetric tree' ([FJ04, p. 47]). This challenge is taken up in the present chapter, in which I answer the following main question.

Main question (Favre and Jonsson)**.** *Is every branchwise-real tree continuously gradable?*

At its heart, this question asks about the existence of a local-global connection for these tree orders. Each branch of a branchwise-real tree X embeds into $\mathbb{R}$, and the question is whether these 'local' embeddings can be combined consistently into one 'global' continuous grading $X \to \mathbb{R}$. In Section 1.3 I will show that the answer to the main question is in fact 'No'. In other words, there is a branchwise-real tree in which these local $\mathbb{R}$-embeddings do not combine into one complete continuous grading.

Theorem 1.7. *There is a branchwise-real tree which is not continuously gradable.*

In order to construct our non-continuously-gradable branchwise-real tree, we will step into the realm of set-theoretic trees (here referred to as 'well-stratified trees', for the sake of clarity). For definitions of the set-theoretic concepts used in remainder of this introduction the reader is referred to Section 1.2. Baumgartner, Laver and Gavin showed in 1970 that there exists a well-stratified tree, all of whose branches are countable, which admits no $\mathbb{R}$-grading (throughout 'countable' means 'finite or countably infinite'). The construction of this tree is given in Section 1.3. Once we have such a tree T, we take what is known as its 'road space': each successor node of T is replaced by a copy of the real interval $(0, 1]$. The resulting tree order is then branchwise-real and has no continuous grading, as required.

The remainder of the chapter is concerned with developing refinements of this technique. In Section 1.4, the link between well-stratified trees and branchwise-real trees is strengthened by proving this following result.

Theorem 1.8. *A branchwise-real tree X is continuously gradable if and only if every well-stratified subtree $T \subseteq X$ is $\mathbb{R}$-gradable.*

The proof works by constructing an increasing sequence $(T_n)_{n\in\omega}$ of 'approximating' well-stratified subtrees of X, whose union is dense in the interval topology; that is, for any $x < y$ in X there is $n \in \omega$ and $z \in T_n$ with $x < z < y$. This technique allows for the application of set-theoretic methods to problems involving branchwise-real trees.

The third part of the chapter focuses on answering certain refinements of the main question, some of which turn out to be independent of ZFC. Can we obtain a non-continuously-gradable branchwise-real tree satisfying certain additional properties? Or, conversely, which properties of branchwise-real trees entail continuous gradability?

Section 1.5 considers a property motivated from the study of $\mathbb{R}$-trees. Those $\mathbb{R}$-trees which appear in applications are often separable (see [Bes01]). The first result is that an $\mathbb{R}$-tree is separable if and only if every cut-point order, regardless of root point, contains at most countably many branching nodes and at most countably many maximal terminal segments isomorphic to a non-trivial real interval. Such orders will be called *countably wispy*. In contrast to the original, the corresponding refinement of the main question has a positive answer, as follows.

Theorem 1.9. *Every countably wispy branchwise-real tree is continuously gradable.*

Generalising, we can ask whether a branchwise-real tree satisfying the natural generalisation of countable wispiness — κ*-wispiness*, for some uncountable cardinal κ — is automatically continuously gradable. Analogously to the countable case, κ-wispiness corresponds to κ-separability on $\mathbb{R}$-trees. In the second part of Section 1.5, I show that the answer to this question in general is independent of ZFC set theory, as follows. Note that part (2) is a generalisation of Theorem 1.9.

Theorem 1.10.

(1) If the Continuum Hypothesis holds then there is an $\aleph_2$-wispy branchwise-real tree with no continuous grading.

(2) If Martin's Axiom holds at κ then all κ-wispy branchwise-real trees are continuously gradable.

In the final section I consider branchwise-real trees satisfying the countable chain condition (ccc). On $\mathbb{R}$-trees, this condition corresponds to a property which is slightly stronger than separability. The answer to the corresponding refinement of the main question is independent of ZFC, as follows.

Theorem 1.11. *The Suslin Hypothesis is equivalent to the statement that every ccc branchwise-real tree has a continuous grading.*

The countable chain condition generalises to the κ chain condition (κ-cc). The correspondent on $\mathbb{R}$-trees is a property slightly stronger than κ-separability. To answer the generalised question 'is every κ-cc branchwise-real tree continuously gradable?', I define the notion of a '$<\kappa$-wide $\mathbb{R}$-ungradable tree', which generalises that of a Suslin tree, and show that the existence of such an object is in general independent of ZFC. Finally, I obtain the following independence result.

Theorem 1.12. *Let κ be an uncountable cardinal. There exists a κ-cc branchwise-real tree with no continuous grading if and only if there exists a $<\kappa$-wide $\mathbb{R}$-ungradable tree.*

The chapter is concluded with a number of open questions.

1.2 Background definitions

The following sets out the main definitions from combinatorial set theory which will come into play in this chapter. For background on these topics, the reader may consult [Jec03, Ch. 9] or [Kun13, §III.3, §II.5]. A *well-stratified tree* T is a tree order in which every $\downarrow(x)$ is well-ordered (in a purely set-theoretic context, we would simply say 'tree'). For any $x \in X$ its *rank* is the order type of $\downarrow(x) \setminus \{x\}$. For any ordinal α, the *αth level* of T, denoted $T(\alpha)$, is the set of elements of rank α. The *height* of T, denoted by $\mathrm{height}(T)$, is the least α such that $T(\alpha)$ is empty. Say that T is *Hausdorff* if it is a meet-semilattice (i.e. satisfies (BR2)).

An *antichain* in a tree order X is a subset $A \subseteq X$ such that for any distinct $x, y \in A$ we have $x \nleqslant y$ and $y \nleqslant x$. Say that X has the *countable chain condition* (ccc) if it has no uncountable antichains.

An *Aronszajn tree* is a well-stratified tree of height ω_1, all of whose levels and branches are countable. A *Suslin tree* is a ccc Aronszajn tree. *Suslin's Hypothesis* (SH) is the statement that there are no Suslin trees. It is well-known that SH is independent of ZFC (see for instance [Jec03, p. 239–242]).

A *forcing poset* is a partial order with a maximum element. In this chapter, while we think of trees as growing upwards, forcing posets are thought of as extending downwards, as is standard in set theory. For this reason, the notion of a 'ccc forcing poset' is defined in the opposite way to the notion of a 'ccc tree order'. Let $\mathbb{P}$ be any forcing poset. An *antichain* in $\mathbb{P}$ is a subset $A \subseteq \mathbb{P}$ such that for any distinct $x, y \in A$ there is no $z \in \mathbb{P}$ such that $z \leqslant x, y$. Then $\mathbb{P}$ has the *countable chain condition* (ccc) if it has no uncountable antichains.

A subset $D \subseteq \mathbb{P}$ is *dense* if for any $p \in \mathbb{P}$ there is $q \leqslant p$ such that $q \in D$. A subset $F \subseteq \mathbb{P}$ is a *filter* if the following hold.

(F1) F is non-empty.

(F2) F is upwards-closed: if $p \in F$ and $p \leqslant q$ then $q \in F$.

(F3) For any $p, q \in F$ there is $r \in F$ such that $r \leqslant p, q$.

For $\kappa < \mathfrak{c}$, *Martin's Axiom* for κ, denoted MA_κ, is the statement that for any ccc forcing poset $\mathbb{P}$ and any collection $\mathcal{D}$ of κ-many dense subsets of $\mathbb{P}$, there is a filter $G \subseteq \mathbb{P}$ which intersects every element of $\mathcal{D}$. *Martin's Axiom* is the statement that MA_κ holds for all $\kappa < \mathfrak{c}$. It is well-known that $\mathrm{MA} + \mathfrak{c} = \aleph_\alpha$ is independent of ZFC, for any regular and uncountable $\aleph_\alpha$ (see [Jec03, Theorem 16.13 and Theorem 16.16]).

Let me now indicate the general mathematical conventions which will be followed in this chapter. A partial function between sets X and Y will be denoted using the notation $f\colon X \rightharpoonup Y$. Let P be any partial order. For any $x \in P$ let:

$$\uparrow(x) := \{y \in P \mid y \geqslant x\}$$

If $S \subseteq P$ is any subset, let:

$$\downarrow S := \{y \in P \mid \exists x \in S\colon y \leqslant x\}, \quad \uparrow S := \{y \in P \mid \exists x \in S\colon y \geqslant x\}$$

For any $x < y$ in P, define:

$$[x, y] := \{z \in P \mid x \leqslant z \leqslant y\}$$

Define the other intervals $[x, y)$, $(x, y]$ and (x, y) analogously. Throughout, unless otherwise specified, 'monotonic function' means 'strictly monotonic function'. Tuples will be denoted using angle brackets: $\langle a, b, \ldots \rangle$.

1.3 Answering the main question

In this section, I answer the main question by proving the following result.

Theorem 1.13. *There is a branchwise-real tree which is not continuously gradable.*

Our non-continuously-gradable branchwise-real tree will be constructed by taking the 'road space' of a certain well-stratified tree. This notion of the road space was first introduced by Floyd Burton Jones [Jon65].

Definition 1.14. Let T be a well-stratified tree. The *road space* of T, denoted $\mathrm{Road}(T)$, is the partial order obtained by replacing each node on a successor level with a copy of the real interval $(0, 1]$, and every other node with a copy of the element 1. Formally, we can view $\mathrm{Road}(T)$ as the following suborder of the lexicographic product order $T \times (0, 1]$:

$$\mathrm{Road}(T) = \{\langle x, t\rangle \in T \times (0, 1] \mid t = 1 \text{ or } x \in T(\alpha + 1) \text{ for some } \alpha\}$$

Note that T embeds canonically in $\mathrm{Road}(T)$ via $x \mapsto \langle x, 1\rangle$.

Lemma 1.15. *When T is a Hausdorff well-stratified tree with no uncountable branches its road space is a branchwise-real tree.*

To prove this, we make use of a basic set-theoretic result, which will reappear often enough to warrant a number.

Lemma 1.16. *Every countable ordinal α embeds into $\mathbb{Q}$ in such a way that limits are preserved.*

Proof. A slick proof makes use of the 'forth' part of the classical 'back-and-forth method'. Enumerate $\alpha = \{\beta_n \mid n \in \omega\}$. We then build up an embedding $f\colon \alpha \to \mathbb{Q}$ inductively on the enumeration. Once we have $f\restriction\{\beta_0, \ldots, \beta_{n-1}\}$, since $\mathbb{Q}$ is a dense linear order without endpoints, we can find $f(\beta_n)$ whose relative position with respect to $f(\beta_0), \ldots, f(\beta_{n-1})$ is the same that of β_n with respect to $\beta_0, \ldots, \beta_{n-1}$. To ensure that limits are also preserved in the resulting embedding, it suffices to require in addition that the distance between $f(\beta_n)$ and its immediate successor in $f(\beta_0), \ldots, f(\beta_{n-1})$ with respect to the order on $\mathbb{Q}$ (if this successor exists) is less than $^1/_{n+1}$. □

Proof Lemma 1.15. Since T is a meet semilattice, so is $\mathrm{Road}(T)$, and hence (BR2) is satisfied. As for (BR1), take any branch B in $\mathrm{Road}(T)$ with the aim of showing that it is isomorphic to a real interval. Let B_T be the result of restricting B to the canonical embedded copy of T in $\mathrm{Road}(T)$. Then B_T is a branch in T, and hence by assumption it is isomorphic to a countable ordinal α. Hence by Lemma 1.16 there is a limit-preserving embedding $\alpha \to \mathbb{Q}$, which then extends to an embedding $B \to \mathbb{R}$, whose image is a real interval. □

Note that if $\ell\colon \mathrm{Road}(T) \to \mathbb{R}$ is a continuous grading, then ℓ restricts to an $\mathbb{R}$-grading of T (via the canonical embedding). Thus, with this lemma, to find a branchwise-real tree X with no continuous grading it suffices to find a Hausdorff, $\mathbb{R}$-ungradable well-stratified tree with no uncountable branches. James E. Baumgartner first constructed a well-stratified tree with these properties, using results from Richard Laver and F. Gavin (see [Bau70]). To understand the construction, let us first examine what it means for a well-stratified tree to be $\mathbb{R}$-gradable, and relate this to $\mathbb{Q}$-gradability. First, $\mathbb{Q}$-gradability is equivalent to the well-known notion of 'specialness'.

Definition 1.17. A well-stratified tree is *special* if it is the union of countably many antichains.

Lemma 1.18. *A well-stratified tree T is $\mathbb{Q}$-gradable if and only if it is special.*

Proof. See [Kun13, Lemma III.5.17]. We only need the forwards direction, which works as follows. Let $f\colon T \to \mathbb{Q}$ be a $\mathbb{Q}$-grading. Let $(q_n)_{n\in\omega}$ be an enumeration of $\mathbb{Q}$. Then for each, $n \in \omega$, let $A_n := f^{-1}\{q_n\}$. Each A_n is an antichain and $T = \bigcup_{n\in\omega} A_n$. □

It turns out that $\mathbb{R}$-gradability is equivalent to the $\mathbb{Q}$-gradability of the nodes of the tree with successor rank.

Definition 1.19. Let $T(\mathsf{succ})$ be the suborder of T consisting of the nodes on its successor levels.

Lemma 1.20. *Let T be a well-stratified tree. Then T is $\mathbb{R}$-gradable if and only if $T(\mathsf{succ})$ is $\mathbb{Q}$-gradable.*

This result is to be found in [Bau70, Ch. 4, Theorem 1(b)]. According to Baumgartner, it is due to Gavin (unpublished).

Proof. Assume that $f\colon T \to \mathbb{R}$ is monotonic. For $x \in T(\mathsf{succ})$ with immediate predecessor y, choose $g(x) \in (f(y), f(x)] \cap \mathbb{Q}$. Then $g\colon T(\mathsf{succ}) \to \mathbb{Q}$ is monotonic. Conversely, assume that $T(\mathsf{succ})$ is $\mathbb{Q}$-gradable. By Lemma 1.18 then we have that $T(\mathsf{succ}) = \bigcup_{n\in\omega} A_n$, where each A_n is an antichain. Define:

$$\begin{aligned} f\colon T &\to \mathbb{R} \\ x &\mapsto \sum \left\{ \frac{1}{n^2} \;\middle|\; \exists y \leqslant x\colon y \in A_n \right\} \end{aligned}$$

Then f is monotonic. □

Given any well-stratified tree T, we can obtain it as $T'(\mathsf{succ})$ of some other tree T' as follows.

Definition 1.21. Let T be a well-stratified tree. Let $\mathrm{Pad}(T)$ be the result of adding a new node directly below every node of T lying on either the 0th level or a limit level.

The following properties are immediate, making use of Lemma 1.20.

Lemma 1.22.

(1) $\mathrm{Pad}(T)(\mathsf{succ}) = T$.

(2) $\mathrm{Pad}(T)$ *is $\mathbb{R}$-gradable if and only if T is $\mathbb{Q}$-gradable.*

(3) T has no uncountable branches if and only if $\mathrm{Pad}(T)$ *has no uncountable branches.*

(4) T is Hausdorff if and only if $\mathrm{Pad}(T)$ *is Hausdorff.*

The task is thus to construct a Hausdorff, $\mathbb{Q}$-ungradable well-stratified tree with no uncountable branches. This is achieved by the following definition and result, due to Laver (unpublished; reported in [Bau70, Ch. 4, Theorem 4(a)]).

Definition 1.23. Let In_ω be the tree of all injective functions of the form $f\colon \alpha \to \omega$, where α is an ordinal, ordered by $\subseteq$. That is:

$$\mathrm{In}_\omega = \{f\colon \alpha \to \omega \mid \alpha \text{ is an ordinal and } f \text{ is injective}\}$$

Theorem 1.24. In_ω *has no* $\mathbb{Q}$*-grading.*

Proof. Following Lemma 1.18, assume for a contradiction that $\mathrm{In}_\omega = \bigcup_{n\in\omega\setminus\{0\}} A_n$, where each A_n is an antichain. We will construct by induction a sequence of elements $f_0 \subset f_1 \subset \cdots$ of In_ω, all with coinfinite range, together with a sequence of natural numbers $x_1, x_2, \ldots$ such that $\mathrm{ran}(f_n) \cap \{x_1, \ldots, x_n\} = \varnothing$. Each x_i represents a 'promise' that it will never appear in the range of an f_n; together they ensure that $\bigcup_{n\in\omega} f_n$ has coinfinite range.

Start with $f_0 := \varnothing$. Assume that f_{n-1} is constructed. Choose any $f_n \in \mathrm{In}_\omega$ and x_n subject to the following conditions.

(i) f_n is a proper extension of f_{n-1}.

(ii) f_n has coinfinite range.

(iii) $\mathrm{ran}(f_n) \cap \{x_1, \ldots, x_n\} = \varnothing$.

(iv) Choose $f_n \in A_n$ if this is possible for some function and x_n satisfying conditions (i), (ii) and (iii).

Note that (i), (ii) and (iii) can be satisfied since f_{n-1} has coinfinite range.

Now let $f := \bigcup_{n\in\omega} f_n$. Then $f \in \mathrm{In}_\omega$, so $f \in A_n$ for some $n > 0$. But note that $\mathrm{ran}(f) \cap \{x_1, x_2, \ldots\} = \varnothing$, so f satisfies (i), (ii) and (iii) above at stage n. Hence by (iv) we must have that $f_n \in A_n$. Then $f_n \subset f$ contradicts that A_n is an antichain. $\lightning$ $\square$

Putting it all together, we can prove Theorem 1.13, thus answering the main question in the negative.

Proof of Theorem 1.13. Take the tree $\mathrm{Pad}(\mathrm{In}_\omega)$. By Theorem 1.24 and Lemma 1.22, this is an $\mathbb{R}$-ungradable well-stratified tree with no uncountable branches. To see that it is Hausdorff, by Lemma 1.22(4) it suffices to show that In_ω is Hausdorff. Take $f, g \in \mathrm{In}_\omega$ distinct. Let α be the least ordinal at which f and g disagree. Then $f{\restriction}\alpha$ is the meet of f and g.

Therefore, by Lemma 1.15, the road space, $\mathrm{Road}(\mathrm{Pad}(\mathrm{In}_\omega))$ is a branchwise-real tree. Furthermore, it has no continuous grading, since any such grading would restrict to an $\mathbb{R}$-grading on $\mathrm{Pad}(\mathrm{In}_\omega)$, via the canonical embedding. $\square$

Theorem 1.25 below shows that the use of an $\mathbb{R}$-ungradable well-stratified tree with no uncountable branches is in some sense essential: every branchwise-real tree with no continuous grading contains such a well-stratified tree.

1.4 The connection to well-stratified trees

In this section, I deepen the connection between branchwise-real trees and well-stratified trees with the following theorem.

Theorem 1.25. *A branchwise-real tree X is continuously gradable if and only if every well-stratified subtree $T \subseteq X$ is $\mathbb{R}$-gradable.*

The proof of the non-trivial direction works by producing a sequence of well-stratified subtrees approximating X. The following is the main construction.

Construction 1.26. Let X be a branchwise-real tree with root p. Fix a well-ordering $(B_\alpha \mid \alpha < \kappa)$ of the branches of X. For each α let $F_\alpha := B_\alpha \setminus \bigcup_{\beta<\alpha} B_\beta$ be the final segment of B_α disjoint from the previous branches. Since X is a branchwise-real tree, each F_α is isomorphic to a real interval (which may be empty or a singleton). Fix an isomorphism $r_\alpha \colon F_\alpha \to I_\alpha$, such that I_α is either empty, the singleton $\{1\}$ or a unit interval. Fix an enumeration $(q_n \mid n \in \omega)$ of $\mathbb{Q} \cap [0, 1]$. Let T_n be the root p of T together with the union of the preimages according to each r_α of the set $\{q_0, \ldots, q_n\}$:

$$T_n := \{p\} \cup \bigcup_{\alpha<\kappa} r_\alpha^{-1}\left(\{q_0, \ldots, q_n\}\right)$$

Lemma 1.27. *Each T_n in Construction 1.26 is a well-stratified tree with no uncountable branches.*

Proof. Note that the segments F_α partition X. Take any $x \in T_n$. Let us see directly that $\downarrow(x)$ is well-ordered. Take any non-empty $S \subseteq \downarrow(x)$. Note that, since S is linearly ordered, if $y \in F_\alpha \cap S$ and $z \in F_\beta \cap S$ with $\alpha < \beta$ then $y < z$. Let $\alpha < \kappa$ be least such that $F_\alpha \cap S \neq \varnothing$. Since $F_\alpha \cap T_n$ is finite, the set $F_\alpha \cap S$ has a least element, which is then least in S.

To see that T_n has no uncountable branches, note that each branch B of T_n is a subset of a branch in X, and thus embeds into $\mathbb{R}$. Since no uncountable ordinal order-embeds into $\mathbb{R}$, we must have that B is countable. □

Because we eventually add every element in the preimages of $\mathbb{Q}$ to the T_n's, the union $\bigcup_{n\in\omega} T_n$ will be a dense subtree of X with respect to the interval topology, as follows.

Definition 1.28. Let X be a tree order with root p. The *interval topology* on X is the topology generated by taking the intervals $[p, x)$ for $x \in X$, (x, y) for $x < y$ in X and $(x, y]$ for $x < y$ in X with y a maximum element, as the basic open sets.[1]

Lemma 1.29. *Let X be a branchwise-real tree and let $(T_n)_{n\in\omega}$ be as in Construction 1.26. Then $\bigcup_{n\in\omega} T_n \subseteq X$ is dense in the interval topology on X. In other words, for any $x < y$ in X there is $n \in \omega$ and $z \in T_n$ with $x < z < y$.*

Proof. Take $x < y$. Let α be least such that $F_\alpha \cap (x, y) \neq \varnothing$. By the minimality of α, the set $F_\alpha \cap (x, y)$ is an initial segment of (x, y), and so is isomorphic to a non-trivial real interval. But then it must contain some $z \in r_\alpha^{-1}(\mathbb{Q})$. This z eventually appears in some T_n, by construction, and satisfies $x < z < y$. $\square$

The last result we need to establish the connection between branchwise-real trees and our increasing sequences of well-stratified trees is that any continuous grading of X corresponds to an $\mathbb{R}$-grading on each of the T_n's. For this, it is first necessary to show that the continuous gradability of X is equivalent to the apparently weaker notion of simple $\mathbb{R}$-gradability. That is, any $\mathbb{R}$-grading of a branchwise-real tree can be transformed into a continuous grading. This is done by 'removing all the gaps', as follows.

Theorem 1.30. *Let X be a branchwise-real tree. Then X has a continuous grading if and only if it has an $\mathbb{R}$-grading.*

Proof. The left-to-right is immediate. So assume that $f\colon X \to \mathbb{R}$ is an $\mathbb{R}$-grading. We will go through the tree eliminating all the discontinuities in f, in a Zorn's Lemma style argument. We work with the set of partial monotonic functions $\ell\colon X \rightharpoonup \mathbb{R}$ such that: (a) $\mathrm{dom}(\ell)$ is downwards-closed, (b) ℓ is continuous on its domain, in the sense of Definition 1.5, and (c) $\ell \leqslant f$ on its domain. By Zorn's Lemma, there is a maximal such partial function $\ell\colon X \rightharpoonup \mathbb{R}$. Suppose for a contradiction that the domain of ℓ is not X.

Pick some maximal linearly ordered subset $C \subseteq X \setminus \mathrm{dom}(\ell)$. If C consists of a single point x, we can extend ℓ to $\widehat{\ell}\colon\ \mathrm{dom}(\ell) \cup \{x\} \to \mathbb{R}$ by setting $\widehat{\ell}(x) := \sup_{y<x} \ell(y)$, noting that this is bounded by $f(x)$. Therefore, we may assume that C is not a singleton. Now, since $\mathrm{dom}(\ell)$ is downwards-closed, C is a final segment in some branch B of X. Hence there is an isomorphism $r\colon C \to I$ onto a unit interval. The map $f \circ r^{-1}\colon I \to \mathbb{R}$ is then monotonic. By a result coming from real analysis, the only

[1]Recall that (x, y) is the interval $\{z \in X \mid x < z < y\}$, etc.

discontinuities on $f \circ r^{-1}$ are jump discontinuities, and there are at most countably many such (see for instance Theorem 4.30 of [Rud76]). We will find some continuous monotonic function lying below $f \circ r^{-1}$. See Figure 1.3 for a picture.

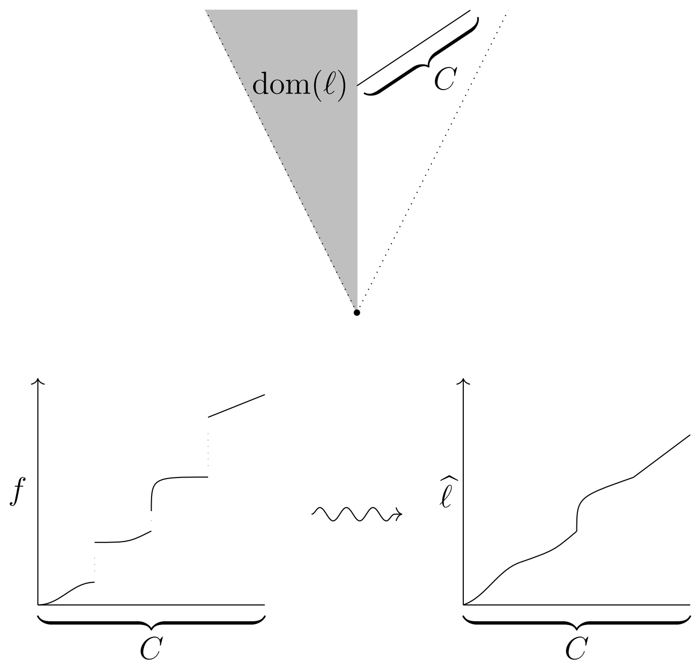

Figure 1.3: Extending ℓ to $\widehat{\ell}\colon \operatorname{dom}(\ell) \cup C \to \mathbb{R}$ by modifying f on C

For $x \in C$, in analogy with real analysis define $f(x-) := \sup_{y<x} f(y)$. Further, define the *jump at x on C* as:

$$j_C(x) := \inf_{\substack{y \in C \\ y > x}} f(y) - f(x-)$$

By the above, there are at most countably many x's on C such that $j_C(x) > 0$. To extend ℓ to C, we first remove each such discontinuity, defining $g_C\colon C \to \mathbb{R}$ by:

$$g_C(x) := f(x-) - \sum_{\substack{y \in C \\ y < x}} j_C(y)$$

Then g_C is continuous with respect to r (i.e. $g_C \circ r^{-1}\colon I \to \mathbb{R}$ is continuous). It is also weakly monotonic, but could fail to be strictly monotonic, e.g. in the case where the jump discontinuities are dense. To obtain a strictly monotonic function, we add a contribution from each jump in a continuous way. Define $h_C\colon C \to \mathbb{R}$ by:

$$h_C(x) := g_C(x) + \sum_{\substack{y \in C \\ y < x}} \frac{r(x) - r(y)}{1 - r(y)} j_C(y)$$

Then function $h_C \circ r^{-1}$ is then the uniform limit of continuous functions, and therefore continuous. Note also that $h_C \leqslant f$ on C.

Finally, it remains to extend ℓ to C by attaching h_C. For this we may need to shift h_C a little so that it fits in continuously with ℓ. Extend ℓ to $\widehat{\ell}\colon \mathrm{dom}(\ell) \cup C \to \mathbb{R}$ by letting, for $x \in C$:

$$\widehat{\ell}(x) := h_C(x) + \left(\sup_{z \in B \setminus C} \ell(z) - \inf_{y \in C} h_C(y) \right)$$

The resulting function is then satisfies the continuity condition, and is moreover such that $\widehat{\ell} \leqslant f$ on its domain. This contradicts the maximality of ℓ. ↯ ☐

Finally, we can establish the connection between the continuous gradability of X and the $\mathbb{R}$-gradability of each T_n in Construction 1.26.

Theorem 1.31. *Let X be a branchwise-real tree and let $(T_n)_{n\in\omega}$ be as in Construction 1.26. Then X is continuously gradable if and only if every T_n is $\mathbb{R}$-gradable.*

Proof. Firstly, any continuous grading of X restricts to an $\mathbb{R}$-grading on each T_n. Conversely, assume that each T_n has an $\mathbb{R}$-grading $f_n\colon T_n \to [0, 1)$. Now, each function f_n can be extended to a weakly monotonic function $\widehat{f}_n\colon X \to [0, 1]$ by:

$$\widehat{f}_n(x) := \sup\{f_n(y) \mid y \in T_n \text{ and } y \leqslant x\}$$

Then define $f\colon X \to [0, 1]$ by:

$$f(x) := \sum_{n\in\omega}^{\infty} \frac{\widehat{f}_n(x)}{2^n}$$

Let us see that f it is strictly monotonic. Take $x < y$ in X. By Lemma 1.29 there is $n \in \omega$ and $x_0, y_0 \in T_n$ such that $x < x_0 < y_0 < y$. Since f_n is monotonic on T_n, we have $\widehat{f}_n(x_0) < \widehat{f}_n(y_0)$. Since also $\widehat{f}_m(x_0) \leqslant \widehat{f}_m(y_0)$ for all $m \in \omega$, we get that $f(x_0) < f(y_0)$. Hence:

$$f(x) \leqslant f(x_0) < f(y_0) \leqslant f(y)$$

Therefore $f\colon X \to \mathbb{R}$ is an $\mathbb{R}$-grading, and thus, by Theorem 1.30, X is continuously gradable. ☐

This last piece allows us to finish the proof of this section's main result.

Proof of Theorem 1.25. If $\ell\colon X \to \mathbb{R}$ is a continuous grading, then it restricts to an $\mathbb{R}$-grading of any well-stratified subtree. Conversely, assume that X has no continuous grading. Let $(T_n)_{n\in\omega}$ be as in Construction 1.26. Then by Theorem 1.31, at least one T_n must be $\mathbb{R}$-ungradable (in fact, infinitely many are). ☐

1.5 Separability and wispiness

With the result of Section 1.4 established, we can now move on to consider refinements of the main question. Our original question was motivated by considering the underlying orders of $\mathbb{R}$-trees. What happens if we look instead at *separable* $\mathbb{R}$-trees? Lemma 1.36 below shows that an $\mathbb{R}$-tree is separable if and only if any cut-point order is what I call 'countably wispy': that it contains fewer than countably many branching nodes, and fewer than countably many terminal segments isomorphic to an non-trivial interval. The corresponding refinement of the main question then becomes the following. In contrast to the main question, the answer turns out to be 'Yes'.

Question 1.32. *Is every countably wispy branchwise-real tree continuously gradable?*

The first item of business is to make the term 'countably wispy' precise, and establish its correspondence with separability.

Definition 1.33. Let X be a tree order and $x \in X$. A *connected component above* x is an equivalence class of $\uparrow(x) \setminus \{x\}$ under the relation:

$$y \sim_x z \quad \Leftrightarrow \quad \text{there is } w > x \text{ such that } w \leqslant y, z$$

The *degree*, $\deg(x)$, of x is the number of connected components above x. Say that x is *terminal* if $\deg(x) = 0$. Say that x is *branching* if $\deg(x) > 1$. Define $\mathrm{Branch}(X)$ to be the set of all branching nodes in X.

Note that we can endow any partial order with the so-called 'Alexandrov topology', in which the open sets are exactly those which are upwards-closed. Under this topology, the connected components above x are precisely the maximal connected subsets of $\uparrow(x) \setminus \{x\}$.

Definition 1.34. A *twig* in a branchwise-real tree is a maximal upwards-closed subset isomorphic to a non-trivial real interval.

See Figure 1.4 for a picture of a twig.

Definition 1.35. A branchwise-real tree is a *countably wispy* if it contains at most countably many branching nodes and at most countably many twigs.

Note that a countably wispy tree may still contain uncountably many terminal nodes. For example, the road space of the well-stratified tree $2^{\leqslant\omega}$ has continuum-many terminal nodes, but no twigs and only countably many branching nodes.

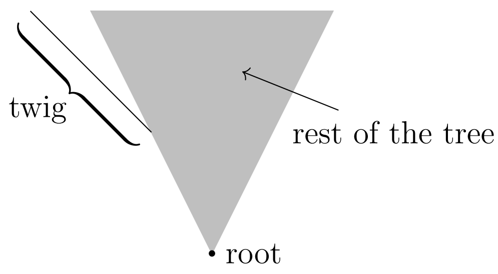

Figure 1.4: What a twig looks like

Lemma 1.36. *An $\mathbb{R}$-tree $\langle X, d\rangle$ is separable if and only if each cut-point order is countably wispy.*

Proof. Fix a cut-point order on X with root p. Assume $D \subseteq X$ is countable and dense in the metric topology. First note that for any $x \in \mathrm{Branch}(X)$, any connected component C above x and any $y \in C$ the open ball about y of radius $d(x, y)$ is open and contained in C; hence $C \cap D \neq \varnothing$. In particular, every branching node sits below an element of D. Now, take any $y \in D$. For every branching node $x < y$, by the above method we can pick $x^* \in D$ which is not in the same connected component above x as y. All the x^*'s must be distinct, and hence the number of branching nodes below y must be countable. Thus, as the countable union of countable sets is countable, there can only be countably many branching nodes in total. Finally, note that any two distinct twigs are disjoint, and that any twig must intersect with D; hence there are only countably many twigs.

Conversely, assume that the cut-point order with root p is countably wispy. Let X_0 be the set of terminal nodes in X, and let X_{t} be the union of the twigs in X. For any $y \in X \setminus (X_0 \cup X_{\mathrm{t}})$, there is $x \in \mathrm{Branch}(X)$ such that $x \geqslant y$. Pick for each $x \in \mathrm{Branch}(X)$ a branch B such that $x \in B$. Furthermore, for each twig there is a unique branch which contains it. Thus we can find a countable set $\mathcal{B}$ of branches through X such that $X = \bigcup \mathcal{B} \cup X_0$. In particular, $\bigcup \mathcal{B}$ is dense in X. Therefore the set:

$$D := \bigcup_{B \in \mathcal{B}} \{y \in B \mid d(p, y) \in \mathbb{Q}\}$$

is countable and dense in X. □

With the connection thus established, we can now answer Question 1.32. In fact, we only need half of the definition of 'countably wispy'.

Theorem 1.37. *Any branchwise-real tree containing at most countably many branching nodes is continuously gradable.*

The proof of this result goes via Theorem 1.25: to show that a branchwise-real tree is continuously gradable, it suffices to show than any well-stratified subtree is $\mathbb{R}$-gradable. For this, we will make use of the following three general lemmas.

Lemma 1.38. *If T is a well-stratified tree with no uncountable branches and* $\mathrm{Branch}(T)$ *is $\mathbb{Q}$-gradable, then T is $\mathbb{R}$-gradable.*

Proof. By Lemma 1.18 we have $\mathrm{Branch}(T) = \bigcup_{n \in \omega} A_n$, where each A_n is an antichain. Define $f_0 \colon \mathrm{Branch}(T) \to \mathbb{R}$ by:

$$f_0(x) := \sum \left\{ \frac{1}{2^n} \;\middle|\; \exists y \leqslant x \colon y \in A_n \right\}$$

Note that f_0 is a bounded $\mathbb{R}$-grading on $\mathrm{Branch}(T)$ with the property that for any $x \in \mathrm{Branch}(T)$:

$$f_0(x) > \sup_{y<x} f_0(y)$$

In other words, below each x there is a 'jump' of size $f_0(x) - \sup_{y<x} f_0(y)$. We now extend f_0 to $f \colon T \to \mathbb{R}$. Define an equivalence relation on $T \setminus \mathrm{Branch}(T)$, setting $u \sim v$ if and only if either $u \leqslant v$ or $v \leqslant u$, and:

$$\downarrow(u) \cap \mathrm{Branch}(T) = \downarrow(v) \cap \mathrm{Branch}(T)$$

Since T has no uncountable branches, each equivalence class is countable; it is moreover well-ordered, and so order-isomorphic to a countable ordinal. Let C be any such class. If there is $x \in \mathrm{Branch}(T)$ such that $x > u$ for all $u \in C$, then there is a least such x. By Lemma 1.16 we can define f monotonic on C such that for all $u \in C$:

$$\sup_{y<x} f_0(y) < f(u) < f_0(x)$$

If there is no such x, then as f_0 is bounded, we can again define f monotonic on C such that for all $u \in C$:

$$\sup_{y<x} f_0(y) < f(u)$$

Putting it all together, we arrive are our desired $\mathbb{R}$-grading $f \colon T \to \mathbb{R}$. □

Lemma 1.39. *If X is a branchwise-real tree, then X is a complete meet-semilattice: any non-empty $S \subseteq X$ has a greatest lower bound $\bigwedge S$.*

Proof. Take any $x \in S$ and consider the set $R := \{x \wedge y \mid y \in S\}$. Since $\downarrow(x)$ is linearly ordered, so is R. Hence R lies in some branch B of X. Since B is order-isomorphic to a real interval, R has an infimum in B, which infimum is then also the element $\bigwedge S$. □

Lemma 1.40. *Let X be a branchwise-real tree and let $T \subseteq X$ be a well-stratified subtree. Then* $|\mathrm{Branch}(T)| \leqslant |\mathrm{Branch}(X)|$.

Proof. Define a function $s\colon \mathrm{Branch}(T) \to \mathrm{Branch}(X)$ as follows. For any $x \in \mathrm{Branch}(T)$, let S_x be its set of immediate successors in T (of which there are at least 2). Then let $s(x) := \bigwedge S_x$ in X. As $|S_x| \geqslant 2$, we have $s(x) \in \mathrm{Branch}(X)$. Let us see that s is injective. Assume that $s(x) = s(y)$ but $x \neq y$. Since $x, y \leqslant s(x)$ in X, which is a tree order, we must have that x and y are comparable, say $x < y$. But then x has an immediate successor z in T such that $z \leqslant y$, and $s(x) < z$ (since x has more than one immediate successor). This means that:

$$s(x) < z \leqslant y \leqslant s(y)$$

which is a contradiction. ↯ □

We can now prove the theorem, answering Question 1.32.

Proof of Theorem 1.37. Let X be branchwise-real tree containing at most countably many branching nodes. Take any $T \subseteq X$ a well-stratified tree. By Lemma 1.40 we have that $|\mathrm{Branch}(T)| \leqslant |\mathrm{Branch}(X)| \leqslant \aleph_0$. So $\mathrm{Branch}(T)$ is a countable well-stratified tree. This means that it has countable height α. By Lemma 1.16, there is a monotonic map $\alpha \to \mathbb{Q}$, which then pulls back through the rank function to a $\mathbb{Q}$-grading on $\mathrm{Branch}(T)$. Then by Lemma 1.38, the tree T is $\mathbb{R}$-gradable. Thus, any well-stratified subtree of X is $\mathbb{R}$-gradable, and so by Theorem 1.25 we get that X is continuously gradable. □

The notion of countable-wispiness generalises naturally to that of κ-wispiness, about which we can ask analogous questions to the above.

Definition 1.41. A branchwise-real tree is a *κ-wispy* if it contains fewer than κ-many branching nodes and fewer than κ-many twigs.

On the metric space side, for uncountable κ this property corresponds to the following generalisation of separability.

Definition 1.42. Let κ be an infinite cardinal. A topological space is *κ-separable* if it admits a dense subset of size less than κ.

Lemma 1.43. *Let κ be an infinite cardinal. An $\mathbb{R}$-tree $\langle X, d\rangle$ is κ-separable if and only if each cut-point order is κ-wispy.*

Proof. The only trouble encountered when generalising the proof of Lemma 1.36 is that, in showing that separability implies that there are only countably many branching nodes, we used that the countable union of countable sets is countable. So we need to be a little careful when κ is not regular. But note that if X is κ-separable for κ a limit cardinal, then it is λ-separable for some regular $\lambda < \kappa$. $\square$

Now, in contrast with the countable case, the existence of an $\aleph_\alpha$-wispy branchwise-real tree with no continuous grading, for $\alpha \geqslant 2$, is independent of ZFC.

Theorem 1.44. *Let* $X = \mathrm{Road}(\mathrm{Pad}(\mathrm{In}_\omega))$ *be the tree constructed in Section 1.3. Then* X *is* $\mathfrak{c}^+$*-wispy. Hence under the Continuum Hypothesis there is a* $\aleph_2$*-wispy branchwise-real tree with no continuous grading.*

Proof. The structure $\mathrm{Branch}(X)$ is isomorphic to $\mathrm{Branch}(\mathrm{In}_\omega)$. An element of In_ω is an injective function from a countable ordinal into ω. There are $\aleph_1$-many countable ordinals, and each has at most $\mathfrak{c}$-many injective functions into ω. Hence:

$$|\mathrm{Branch}(X)| = |\mathrm{Branch}(\mathrm{In}_\omega)| \leqslant |\mathrm{In}_\omega| = \aleph_1 \cdot \mathfrak{c} = \mathfrak{c}$$

Furthermore, a twig in X corresponds to a terminal node in In_ω, of which there are $\mathfrak{c}$-many. Thus X is $\mathfrak{c}^+$-wispy. $\square$

Theorem 1.45. *Let* κ *be a cardinal. Under* MA_κ *every branchwise-real tree* X *with at most* κ*-many branching nodes is continuously gradable.*

Proof. Note that when $\kappa \leqslant \aleph_0$ this already follows from Theorem 1.37. Take any $T \subseteq X$ a well-stratified subtree. By Theorem 1.25, it suffices to show that T is $\mathbb{R}$-gradable. By Lemma 1.38, it suffices to find a $\mathbb{Q}$-grading of $\mathrm{Branch}(T)$. By Lemma 1.40, we have that $|\mathrm{Branch}(T)| \leqslant |\mathrm{Branch}(X)| \leqslant \kappa$. We find a $\mathbb{Q}$-grading for $\mathrm{Branch}(T)$ using the usual 'specialising forcing'.

Let $\mathbb{P}$ be the partial order consisting of all finite, monotonic partial functions $p \colon \mathrm{Branch}(T) \rightharpoonup \mathbb{Q}$ under:

$$p \leqslant q \quad \Leftrightarrow \quad p \text{ extends } q$$

It is well-known that, as $\mathrm{Branch}(T)$ has no uncountable branches, this poset is ccc (see for example Theorem III.5.19 in [Kun13]). Furthermore, for any $x \in \mathrm{Branch}(T)$ the following set is dense in $\mathbb{P}$.

$$D_x := \{p \in \mathbb{P} \mid x \in \mathrm{dom}(p)\}$$

By MA_κ then, using that $|\mathrm{Branch}(T)| \leqslant \kappa$, there is a filter G on $\mathbb{P}$ which intersects with every D_x. Letting $f := \bigcup G$, we see that it is a monotonic function $\mathrm{Branch}(T) \to \mathbb{Q}$; in other words, a $\mathbb{Q}$-grading. □

Remark 1.46. An alternative way of showing that X is continuously gradable is as follows. Enumerate the branches of X as $B_\alpha \mid \alpha < \lambda$, and let $F_\alpha := B_\alpha \setminus \bigcup_{\beta<\alpha} B_\beta$. For each α fix an order isomorphism $r_\alpha \colon F_\alpha \to I_\alpha$ onto a real interval. Then let:

$$X_\mathbb{Q} := \bigcup_{\alpha<\lambda} r_\alpha^{-1}(\mathbb{Q})$$

This is a tree order in which every branch is countable. We can then take the specialising forcing of finite, monotonic partial functions $X_\mathbb{Q} \rightharpoonup \mathbb{Q}$. It is possible to show that this forcing poset is ccc (e.g. by adapting the proof of Theorem III.5.19 in [Kun13]), and hence by a density argument MA_κ entails that there is a $\mathbb{Q}$-grading $X_\mathbb{Q} \to \mathbb{Q}$. This can then be completed to an $\mathbb{R}$-grading of X, which by Theorem 1.30 means that X must be continuously gradable.

Corollary 1.47. *Assume that* ZFC *is consistent. For $\alpha \geqslant 2$ it is consistent that there exists an $\aleph_\alpha$-wispy branchwise-real tree with no continuous grading, and it is also consistent that such a tree does not exist.*

Proof. It is a classical result that the Continuum Hypothesis is consistent with ZFC (see for instance Theorem 13.20 in [Jec03]). Hence by Theorem 1.44 it is consistent that there exists a $\aleph_2$-wispy branchwise-real tree with no continuous grading. Such a tree is also κ-wispy for any $\kappa \geqslant \aleph_2$. Conversely, $\mathrm{MA}_{\aleph_\alpha}$ is consistent with ZFC, and hence by Theorem 1.45 so is the statement that all $\aleph_\alpha$-wispy branchwise-real trees are continuously gradable. □

1.6 The countable chain condition

When building our non-continuously-gradable branchwise-real tree in Section 1.3, we looked for an $\mathbb{R}$-ungradable well-stratified tree with no uncountable branches. There we constructed one explicitly in ZFC, but an observant reader may have noticed that a Suslin tree also satisfies these requirements. Motivated by this observation, in this section we consider branchwise-real trees satisfying the countable chain condition and its generalisations. The principle question is the following refinement of the main question.

Question 1.48. *Is every ccc branchwise-real tree continuously gradable?*

Remark 1.49. Note that in the well-stratified case, ccc-ness plus $\mathbb{R}$-gradability implies that the tree has countable height. Indeed, any ccc, $\mathbb{R}$-gradable well-stratified tree of height ω_1 must be a Suslin tree, but no Suslin tree admits an $\mathbb{R}$-grading (see [Kun13, Lemma IV.6.5]).

Before delving into the details, let us first see what property of $\mathbb{R}$-trees corresponds to ccc-ness.

Definition 1.50. Let X be a connected topological space. A *noncut-point* in X is an element $x \in X$ such that $X \setminus \{x\}$ is connected.

Theorem 1.51. *Let X be an $\mathbb{R}$-tree. The following are equivalent.*

(1) X is separable and has only countably many noncut-points.

(2) Every cut-point order on X is countably wispy and has only countably many terminal nodes.

(3) Every cut-point order on X is ccc.

For this we make use of the following lemma due to Bowditch, which shows that on $\mathbb{R}$-trees connectedness coincides with the natural notion of convexity.

Lemma 1.52. *A subset $Y \subseteq X$ of an $\mathbb{R}$-tree is connected if and only if $[x, y] \subseteq Y$ for every $x, y \in Y$.*[2]

Proof. This is Lemma 1.4 of [Bow99]. □

Proof of Theorem 1.51. Throughout, we fix $p \in X$ and consider the cut-point order on X with root p.

(1) ⇔ (2). By Lemma 1.36, X is separable if and only if every cut-point order is countably wispy. Let X_0 be the set of terminal nodes with respect to the cut-point order. It suffices then to show that the set of noncut-points is X_0 plus possibly p. By Lemma 1.52, it is clear that every element of X_0 is a noncut-point. Conversely, take any noncut-point x in X other than p. If x is not a terminal node, then there is $y \in X$ with $p < x < y$. But then, in $X \setminus \{x\}$, we have p lying in a different connected component to y, since the geodesic $[p, y]$ is not a subset of $X \setminus \{x\}$. ↯

(3) ⇒ (2). That there are only countably many terminal nodes and twigs is immediate. Let $\ell \colon X \to \mathbb{R}$ be the continuous grading on X found using Theorem 1.6. For each $q \in \mathbb{Q}$, the set $\ell^{-1}\{q\}$ is an antichain, which is therefore countable. Below each

[2] Recall that $[x, y]$ denotes the unique geodesic segment between x and y.

such antichain there can only countably many branching nodes, and every branching node must lie below such an antichain. Thus in total there are only countably many branching nodes.

(2) $\Rightarrow$ (3). Take $A \subseteq X$ an antichain. Let A' be the result of removing every element of A which is terminal or lies on a twig. Since each twig can contain at most one element of A, it suffices to show that A' is countable. For each element $x \in A'$, we can pick $x^* \geqslant x$ which is a branching node. For $x \neq y$ we must have $x^* \neq y^*$, since otherwise x and y would be comparable. Since there are only countably many branching nodes, we have that A' is countable. $\square$

Remark 1.53. The equivalence (2) $\Leftrightarrow$ (3) shows that on continuously gradable branchwise-real trees, ccc-ness is only just stronger than countable wispiness. Note however that the proof of the direction (3) $\Rightarrow$ (2) makes key use of the continuous grading. We cannot therefore use Theorem 1.37 to show that every ccc branchwise-real tree is continuously gradable.

We turn to Question 1.48. The following two theorems show that the answer is independent of ZFC. Specifically, there exists a ccc branchwise-real tree with no continuous grading if and only if there exists a Suslin tree.

Theorem 1.54. *The road space of a Suslin tree is a ccc branchwise-real tree with no continuous grading.*

Proof. Let T be a Suslin tree. As noted in Remark 1.49, T admits no $\mathbb{R}$-grading. Therefore $\mathrm{Road}(T)$ has no continuous grading. Furthermore, if $A \subseteq \mathrm{Road}(T)$ is any antichain, then by shifting the elements of A up a little, we may assume that A is contained in the canonical embedded copy of T in $\mathrm{Road}(T)$ (i.e. $T \times \{1\}$). Therefore, A must be countable, and $\mathrm{Road}(T)$ must be ccc. $\square$

Theorem 1.55. *Every ccc branchwise-real tree with no continuous grading contains a Suslin tree.*

Proof. Let X be a ccc branchwise-real tree with no continuous grading. By Theorem 1.25, there is $T \subseteq X$ an $\mathbb{R}$-ungradable well-stratified subtree with no uncountable branches. Since T is a suborder of X, it is also ccc. Furthermore, it must have height ω_1. If not, it would have height α a countable ordinal. Since, by Lemma 1.16, α embeds into $\mathbb{Q}$, this would yield a $\mathbb{Q}$-grading on T. $\lightning$ Thus T is a Suslin tree. $\square$

The countable chain condition generalises to the κ chain condition, which leads to a natural generalisation of Question 1.48.

Definition 1.56. A tree order X is *κ-cc*, where κ is a cardinal, if X has no antichains of size κ.

With a little extra work, we obtain a generalisation of Theorem 1.51 as follows.

Theorem 1.57. *Let X be an $\mathbb{R}$-tree, and let κ be an uncountable cardinal. The following are equivalent.*

(1) X is κ-separable and has fewer than κ-many noncut-points.

(2) Every cut-point order on X is κ-wispy and has fewer than κ-many terminal nodes.

(3) Every cut-point order on X is κ-cc.

Proof. All directions in the proof of Theorem 1.51 readily generalise, except for (3) $\Rightarrow$ (2), in which we use that the union of the countably many branching nodes below each $\ell^{-1}\{q\}$ is countable. We therefore need a new argument in the case where κ is of countable cofinality. I will show that for such a κ, any κ-cc, continuously gradable branchwise-real tree has fewer than κ-many branching nodes.

So, let κ be uncountable with $\mathrm{cf}(\kappa) = \aleph_0$. Then, we can choose a sequence $\kappa_1 < \kappa_2 < \cdots$ of uncountable regular cardinals with limit κ. Let X be a continuously gradable branchwise-real tree which contains at least κ-many branching nodes. We need to find an antichain in X of size κ. This will be constructed in stages $A_0, A_1, \ldots$, all of which are maximal antichains, such that $|A_n| \geqslant \kappa_n$ for $n \geqslant 1$.

Start with A_0 any maximal antichain. Now assume that we have constructed A_n. If $|A_n| \geqslant \kappa$ we can stop the construction here. Otherwise, consider the following partition of X:

$$X = ((\downarrow A_n) \setminus A_n) \cup \bigcup_{x \in A_n} \uparrow(x)$$

The number of branching nodes in $(\downarrow A_n) \setminus A_n$ is at most $|A_n|$. Since the number of branching nodes in X is at least κ, and $|A_n| < \kappa$, there must be $x_n \in A_n$ such that the number of branching nodes in $\uparrow(x_n)$ is at least κ_{n+1}. Now, $\uparrow(x_n)$ is a continuously gradable branchwise-real tree, and so by the generalisation of the proof of (3) $\Rightarrow$ (2) in Theorem 1.51 to the regular cardinal κ_{n+1}, there must be a maximal antichain $A'_{n+1} \subseteq \uparrow(x_n)$ of size at least κ_{n+1}. Let:

$$A_{n+1} := A_n \cup A'_{n+1} \setminus \{x_n\}$$

Finally, let:

$$A := \bigcup_{n\in\omega} A_n \setminus \{x_0, x_1, \dots\}$$

This is then an antichain in X of size κ, as required. □

Remark 1.58. As a corollary of this result, we get that if X is κ-cc for κ some singular cardinal, then it is λ-cc for some $\lambda < \kappa$. This is a specific case of a more general result due to Paul Erdös and Alfred Tarski [ET43].

To show that the existence of a κ-cc branchwise-real tree with no continuous grading is independent of ZFC, we can consider the following notion, which coincides with the notion of Suslin tree for $\kappa = \aleph_1$.

Definition 1.59. Let κ be an uncountable cardinal. A *$<\kappa$-wide $\mathbb{R}$-ungradable tree* is an $\mathbb{R}$-ungradable, κ-cc well-stratified tree with no uncountable branches.

Proposition 1.60.

(1) A Suslin tree is exactly a $<\aleph_1$-wide $\mathbb{R}$-ungradable tree.

(2) MA_κ *implies there are no $<\kappa^+$-wide $\mathbb{R}$-ungradable trees, for $\kappa \geqslant \aleph_1$.*

(3) If $\mathrm{MA} + \neg\mathrm{CH}$ *holds, then there are no $<\mathfrak{c}$-wide $\mathbb{R}$-ungradable trees.*

(4) There is a $<\mathfrak{c}^+$-wide $\mathbb{R}$-ungradable tree.

Proof. (1) This follows since, as noted in of Remark 1.49, all Suslin trees are $\mathbb{R}$-ungradable.

(2) Let T be any κ^+-cc well-stratified tree of height ω_1 with no uncountable branches. We use the specialising forcing from Theorem 1.45 to show that T has a $\mathbb{Q}$-grading (and hence an $\mathbb{R}$-grading). Let $\mathbb{P}$ be the partial order consisting of all finite, monotonic partial functions $p\colon T \rightharpoonup \mathbb{Q}$ under:

$$p \leqslant q \quad\Leftrightarrow\quad p \text{ extends } q$$

As in the proof of Theorem 1.45, $\mathbb{P}$ is ccc, and for each $x \in T$ the following set is dense in $\mathbb{P}$.

$$D_x := \{p \in \mathbb{P} \mid x \in \mathrm{dom}(p)\}$$

To count the size of T, note that each level $T(\alpha)$ has size at most κ, and that there are $\aleph_1$-many levels; hence $|T| \leqslant \kappa \cdot \aleph_1 = \kappa$. Therefore, by MA_κ, there is a filter G which intersects D_x for every $x \in T$. Then $\bigcup G$ is a $\mathbb{Q}$-grading $T \to \mathbb{Q}$.

(3) Let T be any $\mathfrak{c}$-cc well-stratified tree of height ω_1 with no uncountable branches. Then T is the $\aleph_1$-union of its levels, each of which have size less than $\mathfrak{c}$. Now, MA implies that $\mathfrak{c}$ is regular (see e.g. [Kun13, p. 176]). In particular it has cofinality strictly greater than $\aleph_1$, and so the size of T must be less than $\mathfrak{c}$. Since MA holds, by using the specialising forcing on T we obtain a $\mathbb{Q}$-grading $T \to \mathbb{Q}$.

(4) Let $T = \mathrm{Pad}(\mathrm{In}_\omega)$ be the well-stratified tree constructed in Section 1.3. By the proof of Theorem 1.13, T is $\mathbb{R}$-ungradable and has no uncountable branches. As in Theorem 1.44, T has size $\mathfrak{c}$, and hence it is $\mathfrak{c}^+$-cc. □

Corollary 1.61. *Assume that* ZFC *is consistent. For any $\alpha \geqslant 1$, it is consistent that there exists a $<\aleph_\alpha$-wide $\mathbb{R}$-ungradable tree, and it is also consistent that such a tree does not exist.*

Proof. It is consistent that a Suslin tree exists. If this is the case, then by Proposition 1.60 (1) there is a $<\aleph_1$-wide $\mathbb{R}$-ungradable tree, which is also $<\kappa$-wide for all $\kappa \geqslant \aleph_1$. Conversely, it is consistent that $\mathrm{MA}_{\aleph_\alpha}$ holds, and in this case by Proposition 1.60 (2) there are no $<\aleph_\alpha^+$-wide $\mathbb{R}$-ungradable trees. Since any $<\aleph_\alpha$-wide $\mathbb{R}$-ungradable tree is a $<\aleph_\alpha^+$-wide $\mathbb{R}$-ungradable tree, this implies that the former does not exist either. □

Finally, the proofs of Theorems 1.54 and 1.55 readily generalise, yielding the following independence result.

Theorem 1.62. *Let κ be an uncountable cardinal. There exists a κ-cc branchwise-real tree with no continuous grading if and only if there exists a $<\kappa$-wide $\mathbb{R}$-ungradable tree.*

1.7 Open question

Another class of trees finding many applications in geometric group theory is that of Λ-*trees*, a generalisation of $\mathbb{R}$-trees (see [Sha87; Sha91; Mor92; Chi01]). Briefly, an *ordered abelian group* is an abelian group $\langle \Lambda, + \rangle$ equipped with a linear order $\leqslant$ such that if $a \leqslant b$ then $a + c \leqslant b + c$, for every $a, b, c \in \Lambda$. Then, a Λ-*metric space* is a pair $\langle X, d \rangle$ where X is a set and $d \colon X \times X \to \Lambda$ is a function satisfying the usual axioms for metric spaces. We can define intervals in Λ and arcs and geodesic segments in X analogously with the real case. A Λ-*tree* is then a Λ-metric space $\langle X, d \rangle$ such that between any two points $x, y \in X$ there is a unique arc, denoted $[x, y]$, which is also a geodesic segment. When $\Lambda = \mathbb{R}$ we obtain our familiar $\mathbb{R}$-trees, while setting $\Lambda = \mathbb{Z}$ yields the class of graph-theoretic trees.

We can now ask analogues of the main question for ordered abelian groups other than Λ. The notions of branchwise-real order tree and continuous $\mathbb{R}$-grading generalise readily to *branchwise-*Λ *tree order* and *continuous* Λ*-grading*. The main question then generalises to the following classification problem.

Question 1.63. *For which ordered abelian groups* Λ *is it the case that every branchwise-*Λ *tree order is continuously* Λ*-gradable?*

Theorem 1.13 provides a negative answer in the case $\Lambda = \mathbb{R}$. On the other hand, an inductive argument shows that on every branchwise-$\mathbb{Z}$ tree order we can indeed build a continuous $\mathbb{Z}$-grading. What about other ordered abelian groups?

Chapter 2

Uniform, rigid branchwise-real trees

2.1 Introduction

This chapter concerns the construction of branchwise-real trees which are rigid — i.e. without non-trivial order-automorphisms. Branchwise-real trees were first defined in [FJ04].[1] They are tree partial orders in which every branch is order-isomorphic to a real interval, and in which every two elements have a meet (see Section 2.2 for definitions).

Branchwise-real trees were originally motivated from the study of $\mathbb{R}$-trees. Informally, an $\mathbb{R}$-tree is a metric space tree in which every point is permitted to be branching. They play a role in geometric group theory, in which one considers some infinite group acting by isometries on an $\mathbb{R}$-tree. For more information, see [Bes01; MNO92] and the references contained in [Fab15]. Now, fixing any point p in an $\mathbb{R}$-tree, one can consider the *cut-point order*: the set of paths through the tree from p, ordered by path extension. The resulting partial order is a branchwise-real tree. Conversely, any branchwise-real tree which admits a monotonic function into the reals is the cut-point order of some $\mathbb{R}$-tree ([FJ04, p. 50] and Theorem 1.6 above). The present study of the order-automorphisms of branchwise-real trees is thus motivated by the study of the isometries of $\mathbb{R}$-trees.

A different motivation comes the investigation of the automorphisms of well-stratified trees (tree partial orders in which every branch is isomorphic to an ordinal) [GS64; Jen69; Jec72; Avr79; AS85; FH09]. If X is a well-stratified tree without uncountable branches, we can turn it into a branchwise-real tree by taking its 'road space', essentially replacing every node with a copy of the interval $[0, 1)$, so that every

[1] In [FJ04] branchwise-real trees are called 'non-metric trees'.

branch becomes isomorphic to a real interval [Jon65]. Every automorphism of X can then be extended to an automorphism of its road space (but not vice versa).

However, the class of branchwise-real trees is more general than that obtained by taking the road spaces of well-stratified trees without uncountable branches. This is because any point can be branching, so that the structure of the branching points can be far from well-stratified. Consider for instance the 'comb' structure given by taking a 'shaft' $[0, 1]$ and adding a 'tooth' — a copy of $(0, 1]$ — as a new branch from each point on the shaft. See Figure 2.1 for a picture. Every point along the shaft except the maximal one is branching of degree 2. Moreover, the task of constructing a rigid branchwise-real tree is more challenging than in the well-stratified case. Indeed, any part of the tree which looks like a real interval and which contains no branching points allows for an easy automorphism which fixes the rest of the tree and permutes the interval in some non-trivial way.

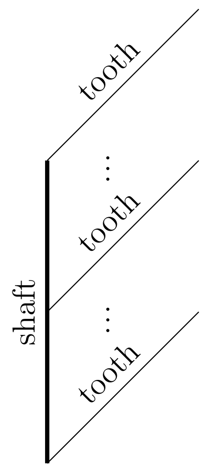

Figure 2.1: A representation of the comb branchwise-real tree, consisting of the shaft and a few example teeth

In this chapter, I construct examples of rigid branchwise-real trees subject to two uniformity conditions, the latter stronger than the former. First, I construct one in which every branching node has the same degree κ, for $2 \leqslant \kappa \leqslant \mathfrak{c}$. I call this condition 'weakly uniformly κ-branching'. Second, using a different technique, I construct one in which every node is branching and of the same degree κ, for $2 \leqslant \kappa \leqslant \mathfrak{c}^+$. This condition I call 'uniformly κ-branching'. This second tree admits a monotonic function into the reals, and thus corresponds to the cut-point order of some $\mathbb{R}$-tree. The technique can be further adapted to produce such a uniformly κ-branching, rigid branchwise real tree which does not admit a monotonic function into the reals, and thus is not related to an $\mathbb{R}$-tree. These results manifest the phenomenon of the marriage of two seemingly opposing notions: rigidity and uniformity. As such they continue in the vein of similar results found for well-stratified trees [Avr79; FH09].

All trees here are grown recursively in stages. We start with a single point: the root. At successor stages, we choose some number of points added at the previous stage, and add $(\kappa - 1)$ many new 'spines' — copies of the interval $(0, \infty)$ — above. (Throughout $\kappa - 1$ denotes $n - 1$ if $\kappa = n$ finite and κ otherwise.) We never extend directly above a spine once it has been added. At limit stages there may be new branches appearing through the tree, consisting of portions added at each of the previous stages. We decide which of these to extend by adding a new point on top.

A common technique used throughout when growing rigid trees is to first produce a collection of colourings of the interval $(0, \infty)$. As the tree grows, we take for each new spine added a different, unused colouring and lay it along the spine. In this way, we also build up a colouring of the whole tree as we go. These colourings are then used to determine which points will form the bases of new spines at successor stages, and which new branches to extend at limit stages. The collection of colourings of $(0, \infty)$ is carefully constructed so as to enable the desired properties of the final tree.

The constructions of the trees in this chapter, though they occur in ZFC, have strong forcing flavour to them. Indeed, I first built these rigid trees using forcing notions, before realising that related constructions could be carried out in ZFC. I will briefly elaborate on two of these forcing notions. For more details on forcing, see [Hal17] and [Jec03; Kun13]. However, knowledge of forcing is not required to understand any of the ZFC constructions in this chapter.

2.2 Background

Let me begin by fixing some terminology and notation relating to partial orders. Let P be a partial order. A *chain* in P is a linearly-ordered subset. A *branch* through P is a maximal chain. A *ray* is a final segment of a branch. A subset $Q \subseteq P$ is *coinitial* if for every $x \in P$ there is $y \in Q$ such that $y \leqslant x$. Elements $x, y \in P$ are *comparable* if $x \leqslant y$ or $y \leqslant x$. Two subsets $Q, R \subseteq P$ are *comparable* if there is $x \in Q$ and $y \in R$ such that x and y are comparable. An *antichain* is a set of pairwise incomparable elements. For any $x \in P$ let:

$$\downarrow(x) := \{y \in X \mid y \leqslant x\}, \qquad \uparrow(x) := \{y \in P \mid y \geqslant x\}$$

For any $x < y$ in P, define:

$$[x, y] := \{z \in P \mid x \leqslant z \leqslant y\}$$

Define the other intervals $[x, y)$, $(x, y]$ and (x, y) analogously. A function $f\colon P \to Q$ between partial orders is *monotonic* if whenever $x < y$ we have $f(x) < f(y)$. I will also call such a function a *Q-grading* of P. An *isomorphism* between P and Q is a bijection $f\colon P \to Q$ which is monotonic with monotonic inverse.

I can now introduce the main object of study: branchwise-real trees. I follow the presentation given in Chapter 1 (see Definition 1.1 and Definition 1.2).

Definition 2.1. A *tree order* is a partial order X such that the following conditions hold.

(TO1) For every $x \in X$ the set $\downarrow(x)$ is a linear order.

(TO2) X has a minimum element, its *root*.

Definition 2.2. A *branchwise-real tree* is a tree order X subject to the following extra conditions.

(BR1) Every branch is order-isomorphic to a real interval.

(BR2) For any $x, y \in X$, the set $\{z \in X \mid z \leqslant x, y\}$ has a maximum element $x \wedge y$, the *meet* of x and y (in other words, X is a meet-semilattice).

We will also meet trees in which every branch is well-ordered. Such trees will be called 'well-stratified trees'; in a purely set theoretic context, we would simply say 'trees'.

Definition 2.3. A tree order T is *well-stratified* if every branch is order-isomorphic to an ordinal. The *rank* of an element $x \in T$ is the order type of $\downarrow(x) \setminus \{x\}$.

In order to define the *degree* of a point in a tree order we need the notion of a connected component above that point.

Definition 2.4. Let X be a tree order and $x \in X$. A *connected component above x* is an equivalence class of $\uparrow(x) \setminus \{x\}$ under the relation:

$$y \sim_x z \quad \Leftrightarrow \quad \text{there is } w > x \text{ such that } w \leqslant y, z$$

Note that, when X is a meet-semilattice, we have $y \sim_x z$ if and only if $y \wedge z > x$. See Figure 2.2 for an example illustration of this relation. I will usually drop the 'connected' and refer to these as 'components above x'. The *degree*, $\deg(x)$, of x is the number of components above x. Say that x is *terminal* if $\deg(x) = 0$. Say that x is *branching* if $\deg(x) > 1$. The *degree* of X is the supremum of the degrees of its elements.

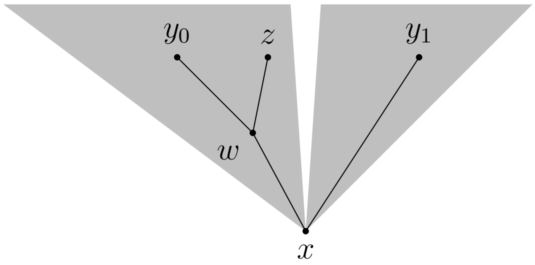

Figure 2.2: An example illustrating the relation $\sim_x$ which defines components above x. We have $y_0 \sim_x z$ because $x < w \leqslant y_0, z$, but $z \not\sim_x y_1$ because there is no element above x which lies below both y_1 and z.

The notion of *continuous gradability* is important for branchwise-real trees and is the main property studied in Chapter 1 (see Definition 1.5).

Definition 2.5. Let X be a branchwise-real tree. An $\mathbb{R}$-grading $\ell \colon X \to \mathbb{R}$ is *continuous* if for any $x < y$ in X the restriction:

$$\ell \restriction [x, y] \colon [x, y] \to [\ell(x), \ell(y)]$$

is an order-isomorphism. I will usually drop the '$\mathbb{R}$' and call such functions 'continuous gradings'. Say that X is *continuously gradable* if it admits a continuous grading.

Theorem 2.6. *A branchwise-real tree is continuously gradable if and only if it is* $\mathbb{R}$*-gradable.*

Proof sketch. See Theorem 1.30. An arbitrary $\mathbb{R}$-grading f of a branchwise-real tree X may contain a number of discontinuities. Using a basic result from real analysis, it can be shown that such discontinuities must take the form of 'jumps' in value, and that any branch may contain only countably many. The proof proceeds to eliminate every jump in f, using a Zorn's Lemma style argument. □

Two final pieces of notation. Throughout, tuples will be denoted using angle brackets: $\langle a, b, \ldots \rangle$. A partial function between sets X and Y will be denoted using $f \colon X \rightharpoonup Y$.

2.3 How to grow branchwise-real trees

The trees in this chapter are grown using an iterative process. We always start with a singleton as the 'root'. At a successor step, above each point introduced in the previous stage we can add any number of new 'spines'. A spine is a copy of the positive real numbers $(0, \infty)$, or the interval $(0, \infty]$, which is 'terminal', in the sense that we will never extend the tree above the end of a spine. In other words, in the final tree each spine will be a ray. Call the former type of spine an *open spine* and the latter a *closed spine*. At a limit step, we first take the union of the previous stages. There is more to do however, since although we do not extend above spines, new branches through the tree appear at the limit, which we may choose to extend or not. A newly appearing branch contains the root, follows a stage-1 spine partway, then branches onto a stage-2 spine, and so on, cofinally in the limit. Such a branch thus consists of a little piece from each previous stage of the construction, with no final piece. If we decide to extend it, we add a new point directly above. This process is illustrated in Figure 2.3.

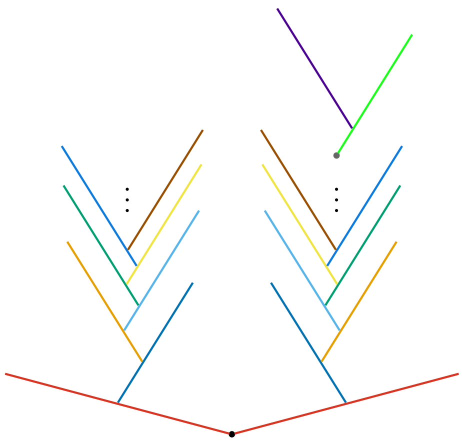

Figure 2.3: An example of growing a branchwise-real tree for $\omega + 2$ steps. Spines beginning at the same vertical height are added at the same stage. Each stage is also represented using a different colour. At stage ω, there are two new branches through the tree, the one on the right and the one on the left. We only extend the right, and we do so by adding a point at the limit, which is then further extended in subsequent stages.

This construction may continue for ω_1-many steps. Note however that since every branch must be order-isomorphic to a real interval, and there is no increasing ω_1-sequence of reals, no branch through the final tree can contain a piece from every stage below ω_1.

Remark 2.7. A similar construction was used by Urysohn in [Ury27] (in German) to construct a metric space which is nowhere separable. Urysohn essentially followed the process outlined above, adding $\mathfrak{c}$-many open spines above each point added at the previous step, for ω-many steps. Of course, Urysohn was constructing a topological space, not a partial order, but there is a clear correspondence between the two types of structure (as elaborated in Chapter 1). In [Ber19], Urysohn's construction is defined (in English) and extended to all countable steps, by taking the Hausdorff completion at limit steps.

Now, given any branchwise-real tree X constructed in the fashion described above, we can naturally define a rank function $\rho\colon X \to \omega_1$, where $\rho(x)$ is the stage α at which x was introduced. As Lemma 2.9 below demonstrates, any branchwise-real tree can be realised using the construction above, and thus admits a rank function. Note that rank functions are not unique in general. In fact, the same tree can have rank functions with a variety of ranges, as Figure 2.4 demonstrates.

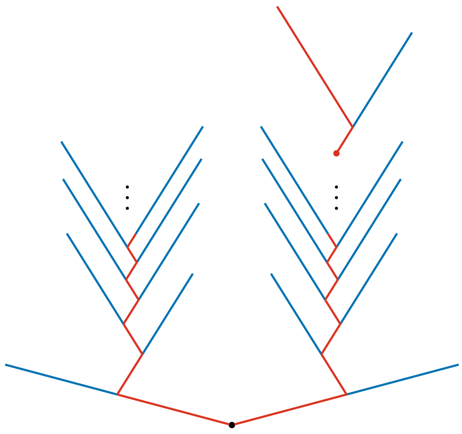

Figure 2.4: An alternative construction of the tree in Figure 2.3, which yields a different rank function. The root gets rank 0. Both the branch which passes through every spine on the left and the branch which passes through every spine on the right have rank 1. The rest of the tree has rank 2. While the original construction produced $\omega + 2$ ranks, this only produces 3.

Definition 2.8. Let X be a tree order. A *rank function* on X is a monotonic function $\rho\colon X \to \gamma$, where γ is some ordinal, such that the following hold.

(RF1) For every branch B in X the set $\{\rho(x) \mid x \in B\}$ is downwards-closed.

(RF2) If α is 0 or a limit, then $\rho^{-1}[\{\alpha\}]$ is an antichain.

(RF3) If α is a successor then $\rho^{-1}[\{\alpha\}]$ is the disjoint union of a family of incomparable rays.

Lemma 2.9. *Every branchwise-real tree admits a rank function.*

Proof. Let X be a tree order. By Zorn's Lemma, there is a maximal partial rank function ρ whose domain is a downwards closed subset $Y \subseteq X$. Suppose for a contradiction that there is $x \in X \setminus Y$. Consider $\downarrow(x) \cap Y$. Note that $\downarrow(x) \cap Y$ is non-empty, since we can always give the root rank 0. Furthermore, we can assign any limit of ranked points the rank which is the limit of the ranks of those points. Hence $\downarrow(x) \cap Y$ has a maximum element y. Let C be the component above y which contains x, and pick a maximal ray R through C which contains x. Then we can extend ρ so that x becomes ranked by giving every element of R rank $\rho(y) + 1$. $\lightning$ □

Remark 2.10. Lemma 2.9 shows that the construction elaborated above of iteratively adding spines is a completely general way of building branchwise-real trees. It was necessary to permit closed spines in addition to open spines in order to allow for trees with terminal nodes. All trees produced in this chapter, however, are without terminal nodes. Hence all constructions from now on will use open spines exclusively.

Let us now establish some facts concerning rank functions on a branchwise-real tree X. Since every branch is order-isomorphic to a real interval, we can assume that the codomain of any rank function on X is ω_1.

Definition 2.11. A rank function $\rho\colon X \to \omega_1$ is *bounded* if the supremum of the set $\{\rho(x) \mid x \in X\}$ is less than ω_1.

Lemma 2.12. *If X has a bounded rank function, then it is continuously gradable.*

To prove this lemma, we make use of the following set-theoretic result due to Cantor [Can95], which in particular shows that every countable ordinal embeds into $\mathbb{Q}$.

Lemma 2.13. *The set $\mathbb{Q}$ of rationals is universal for countable linear orders: every countable linear order X embeds into $\mathbb{Q}$.*

Proof. A slick proof makes use of the 'forth' part of the classical 'back-and-forth method'. Enumerate $X = \{x_n \mid n \in \omega\}$. By induction on this enumeration we then build up an embedding $r\colon X \to \mathbb{Q}$. Once we have $r \restriction \{x_0, \dots, x_{n-1}\}$, since $\mathbb{Q}$ is a dense linear order without endpoints, we can find $r(x_n)$ whose relative position with respect to $r(x_0), \dots, r(x_{n-1})$ is the same that of x_n with respect to $x_0, \dots, x_{n-1}$. □

Proof of Lemma 2.12. Let $\rho\colon X \to \omega_1$ be a rank function with:

$$\gamma := \{\rho(x) \mid x \in X\} < \omega_1$$

Let $r\colon \gamma \to \mathbb{Q}$ be an embedding furnished by Lemma 2.13. We construct an embedding $\ell\colon X \to \mathbb{R}$ by induction on the rank. If α is 0 or a limit, map each element in $\rho^{-1}[\{\alpha\}]$ to $r(\alpha)$. Take a successor $\beta+1$. Then $\rho^{-1}[\{\beta+1\}]$ is the disjoint union of a family of incomparable rays, each of which is order-isomorphic to a real interval. Take for each an isomorphism onto a sub-interval of $(r(\beta), r(\beta+1))$, and use this to extend ℓ there. After γ-many steps, we obtain an $\mathbb{R}$-grading $X \to \mathbb{R}$. Then by Theorem 2.6, we find that X is continuously gradable. □

Remark 2.14. The converse of Lemma 2.12 does not hold. A good counterexample is U_κ, the universal continuously gradable branchwise real tree of degree $\kappa \geqslant 2$, defined as follows. This is essentially the construction of the universal $\mathbb{R}$-tree given in [DP01]; see also [Nik89; MNO92]. Let U_κ be the set of functions $r\colon [0,a) \to \kappa$, for each nonnegative real a, which are 'piecewise constant to the right': for any $t \in [0,a)$ there is $\epsilon > 0$ such that $r{\restriction}[t, t+\epsilon)$ is constant. Note that this condition ensures that the set of 'value-change points' — those $t \in [0,a)$ for which for every $\epsilon > 0$ there is $s \in (t-\epsilon, t)$ such that $r(s) \neq r(t)$ — is well-ordered.

Then U_κ becomes a branchwise-real tree under function extension. We can define a continuous grading ℓ on U_κ by setting $\ell(r) = a$ when $r\colon [0,a) \to \kappa$. With this definition, any continuously graded branchwise-real tree of degree at most κ embeds into U_κ in such a way that the continuous gradings agree.

Let us see briefly why U_κ has no bounded rank function. Suppose $\rho\colon X \to \gamma$ is a (surjective) rank function with $\gamma < \omega_1$. Assume that γ is a limit ordinal, the successor case being similar. By Lemma 2.13 we can find an order embedding $s\colon \gamma+1 \to [0,\infty)$ such that $s(0) = 0$. We can recursively define $r\colon [0, s(\gamma)) \to \kappa$ such that $\rho(r{\restriction}[0, s(\alpha))) = \alpha$ for each $\alpha < \gamma$ as follows. To define r on $[s(\alpha), s(\alpha+1))$, choose a ray above $r{\restriction}[0, s(\alpha))$ of rank $\alpha+1$. This ray is essentially a function $h_\alpha\colon [s(\alpha), \infty) \to \kappa$. Define:

$$r{\restriction}[s(\alpha), s(\alpha+1)) := h_\alpha{\restriction}[s(\alpha), s(\alpha+1))$$

At limits we take unions. But now r cannot have a rank (it would get rank γ), which is a contradiction. ↯

In Section 2.5 below, we shall be considering the interaction between rank functions, connected components and order-automorphisms. The following basic facts will be useful.

Definition 2.15. Let $\langle X, \rho \rangle$ be a ranked tree order. Take $x \in X$ and C a component above x. The *rank* of C is $\rho(C) := \min\{\rho(y) \mid y \in C\}$.

Note that the rank of a connected component is always a successor.

Lemma 2.16. *Let $\langle X, \rho \rangle$ be a ranked tree order and take $x \in X$ non-terminal.*

(1) When $\rho(x)$ is 0 or a limit, every component above x has rank $\rho(x) + 1$.

(2) When $\rho(x)$ is a successor, there is one component above x of rank $\rho(x)$ and the rest are of rank $\rho(x) + 1$.

Let $f : X \to X$ be an order-automorphism.

(3) Then f factors through a bijection of the components above x onto the components above $f(x)$.

(4) If f maps the component C above x onto the component $f(C)$ above $f(x)$, then there is a coinitial interval $I \subseteq C$ of constant rank $\rho(C)$ which maps onto a coinitial interval $f(I) \subseteq f(C)$ of constant rank $\rho(f(C))$.

Proof. (1) Let C be any component above x, and pick $y \in C$ of minimal rank. Then the restriction of ρ to $\downarrow(y)$ is surjective onto $\rho(y) + 1$, from which we conclude that $\rho(C) = \rho(y) = \rho(x) + 1$.

(2) Since $\rho(x)$ is a successor, x lies on a ray all of whose elements have rank $\rho(x)$. Then as ρ is monotonic, the component above x through which this ray passes must have rank $\rho(x)$. Let C be any other component above x. Then C cannot have rank $\rho(x)$, since $\rho^{-1}[\{\rho(x)\}]$ consists of the disjoint union of a family of incomparable rays. By the above argument, we have that $\rho(C) = \rho(x) + 1$.

(3) Let C be any component above x. If $y', z' \in f(C)$, then there is $w > x$ such that $w \leqslant f^{-1}(y'), f^{-1}(z')$. Hence we have $f(x) < f(w) \leqslant y', z'$, so that y' and z' lie in the same component above $f(x)$. By the same argument applied to f^{-1}, we see that f maps components above x onto components above $f(x)$, and vice versa.

(4) There is a ray R through C all of whose elements are of rank $\rho(C)$. Similarly, there is a ray R' through $f(C)$ all of whose elements are of rank $\rho(f(C))$. Take $y' \in R'$ and $z' \in f(R)$. Then there is $w' > f(x)$ such that $w' \leqslant y', z'$. This w' then lies on both rays. Since f is an automorphism, this means that the coinitial interval $(x, f^{-1}(w')]$ is mapped onto the coinitial interval $(f(x), w']$. □

2.4 A weakly uniformly branching, rigid branchwise-real tree

In this section, I construct our first rigid branchwise-real tree. It satisfies the following uniformity condition, the stronger version of which will be considered in the next section.

Definition 2.17. A tree order is *weakly uniformly κ-branching* if every branching node has degree κ.

Theorem 2.18. *Let $2 \leqslant \kappa \leqslant \mathfrak{c}$. There is a weakly uniformly κ-branching, rigid branchwise-real tree which is continuously gradable.*

Proof. We first construct a collection of dense, mutually non-isomorphic sets of positive reals $\{S_A \subseteq (0, \infty) \mid A \subseteq \omega\}$, such that S_A 'encodes' A.[2] The set S_A consists of the rationals in $(0, \infty)$ plus a descending sequence of real intervals with limit 0 whose endpoints are irrational, such that the nth interval is open when $n \notin A$ and half-open when $n \in A$. Formally, fix some descending sequence (x_n) of irrationals with limit 0. Then for any $A \subseteq \omega$ we define:

$$S_A := (\mathbb{Q} \cap (0, \infty)) \cup \bigcup_{n \in \omega} (x_{2n+1}, x_{2n}) \cup \{x_{2n} \mid n \in A\}$$

See Figure 2.5 for a representation of an example S_A.

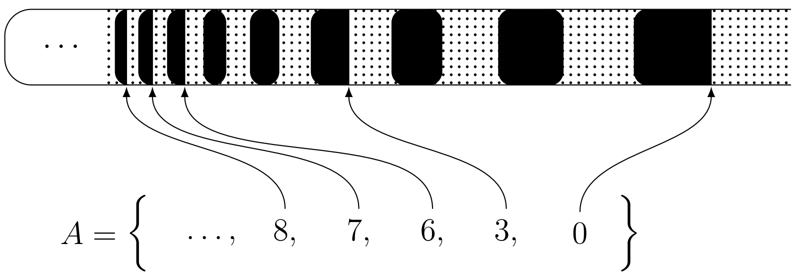

Figure 2.5: A representation of an example S_A set. Solid regions contain every real, while dotted regions only contain the rationals. A curved boundary on a solid region signifies that it does not include that endpoint. The set $A \subseteq \omega$ is written in descending order, so as to coincide with the ordering of the real intervals. Arrows relate the elements of A with the corresponding endpoints of real intervals which code for them in S_A.

[2]In fact, we only need that the S_A's are mutually non-isomorphic as subsets of $(0, \infty)$, but the construction yields sets which are mutually non-isomorphic as stand-alone linear orders.

Suppose that $p\colon S_A \to S_B$ is an order-isomorphism for $A \neq B$. Then p must map the real intervals onto the real intervals. Moreover, the rightmost real interval in S_A must map to the rightmost real interval in S_B, and so on. But A and B disagree on some natural number, say $n \in A \setminus B$ (without loss of generality). Then the nth real interval from the right in S_A is half-open, and p sends it to the nth real interval in S_B, which is open. Since we ensured that the endpoints of the real intervals were irrational, this means that p doesn't preserve the fact that the nth real interval in S_A has a supremum. $\lightning$

Our rigid branchwise-real tree X is now constructed in ω-many stages using open spines, laying along each new spine a different set S_A, and using this to determine which points to use as the bases of new spines. Note that we need to make sure that we have fresh S_A sets available at each stage of the construction: we don't want to run out. For this, we can partition $\mathcal{P}(\omega)$ into ω-many batches of size $\mathfrak{c}$, so that the nth batch is reserved for the nth stage.

We construct X in stages X_n for $n \in \omega$. Start with X_0 the singleton tree, and form X_1 by adding κ-many new open spines above the root. Now assume that X_n is constructed for $n \geqslant 1$. To each spine U added at stage n, associate a different, unused subset $A \subseteq \omega$. Fix an isomorphism $p_U\colon U \to (0, \infty)$ and above each element in $p_U^{-1}(S_A)$ add $(\kappa - 1)$-many new open spines. Finally, let $X := \bigcup_{n\in\omega} X_n$. Define a rank function $\rho\colon X \to \omega$, so that for each $n \in \omega$ we have $\rho^{-1}[\{n+1\}] = X_{n+1} \setminus X_n$.

Note that at each stage of the construction, we made sure that any branching points got degree κ. Furthermore, ρ is a bounded rank function on X, hence by Lemma 2.12 we have that X is continuously gradable.

Let us see that X is rigid. Suppose for a contradiction that $f\colon X \to X$ is a non-trivial order-automorphism. Then there is $x \in X$ such that $f(x) \neq x$. There must then be $y > x$ such that $f([x, y])$ and $[x, y]$ are disjoint. Since each S_A is dense in $(0, \infty)$, there must be $z \in [x, y]$ which is branching. Choose any spine U above z added during the construction. Then $f(U)$ is a ray through $\uparrow(f(z))$ and $f{\restriction}U\colon U \to f(U)$ is an isomorphism such that $w \in U$ is branching if and only if $f(w) \in f(U)$ is branching.

Let us now consider the pattern of branching nodes along U and $f(U)$. Following U, this looks like some S_A: it consists in an alternating sequence of real intervals and rational intervals, which converge at the base. Hence we must have the same pattern along $f(U)$. Now, there three possibilities for the way in which $f(U)$ fits into our construction. (i) $f(U)$ is a single spine added above $f(z)$. (ii) $f(U)$ consists of an initial segment of a spine added above $f(z)$, followed by an initial segment of a spine added at the next stage (followed potentially by further initial segments of spines added at later stages). (iii) $f(U)$ starts part-way up a spine added at an earlier stage. See Figure 2.6 for a representation of the three types of patterns of branching points which can occur along $f(U)$. Considered as subsets of $(0, \infty)$, no two of three possible patterns of branching nodes can be order-isomorphic. Since the pattern of branching nodes along U looks like (i), the pattern along $f(U)$ must also have this form.

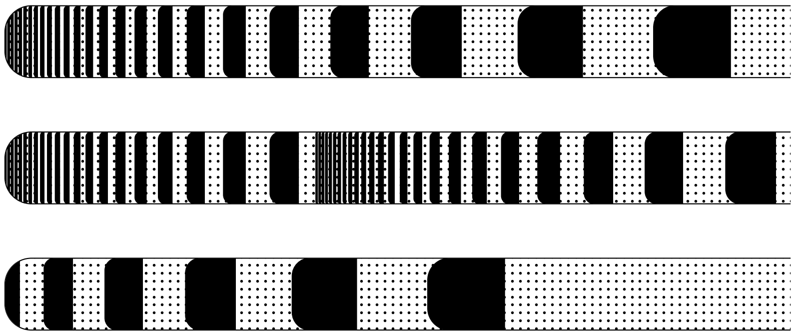

Figure 2.6: A comparison three possible types of pattern which can occur along $f(U)$. Solid regions contain every real, while dotted regions only contain the rationals.

Therefore, there are $A, B \subseteq \omega$ distinct, such that the pattern of branching nodes along U looks like S_A, and the pattern of branching nodes along $f(U)$ looks like S_B. In other words, we have order-isomorphisms $p_U \colon U \to (0, \infty)$ and $p_{f(U)} \colon f(U) \to (0, \infty)$ such that $p_{f(U)} \circ f \circ p_U^{-1} \colon (0, \infty) \to (0, \infty)$ is an order-automorphism mapping S_A onto S_B. This is our contradiction. ↯ □

Let us now consider the relationship between the construction just given and a forcing notion. For ease of exposition, we only deal with the case $\kappa = \aleph_0$. We force a weakly uniformly $\aleph_0$-branching, continuously gradable, rigid branchwise-real tree by considering countable approximations to it. Such objects might be called *sub-branchwise-real trees*: we replace (BR1) with the requirement that each branch be embeddable into $\mathbb{R}$. In addition, we keep track of a rank function from the approximations into ω, together with a set of 'non-extension promises'. The latter is a subset of the sub-branchwise-real tree consisting of degree-1 nodes which guarantees that these nodes never become branching. The forcing conditions are thus triples $\langle X, \rho, S \rangle$ consisting of: (a) a countable sub-branchwise-real tree X (which for concreteness we can take to be a subset of ω_1), (b) a rank function $\rho \colon X \to \omega$, and (c) a non-extension promise set $S \subseteq X$ of degree-1 nodes. For one condition to be stronger than another, we require the sub-branchwise-real tree, rank function and promise set of the former extend those of the latter, and further that we do not add new connected components above any already-branching point, so that in particular the degrees of branching points do not decrease. To ensure this latter condition, we require that for every element x of the weaker approximation, every component above x with respect to the stronger approximation contain at most one component above x with respect to the weaker. Note that the fact that stronger conditions must extend the promise sets and that every element of the promise set must have degree 1 ensures that we keep our promises. Let $\mathbb{P}$ be the forcing notion just described.

Theorem 2.19. *The forcing notion $\mathbb{P}$ is countably closed, and any $\mathbb{P}$-generic set determines an ω-ranked, weakly uniformly $\aleph_0$-branching, rigid branchwise-real tree.*

Proof sketch. That $\mathbb{P}$ is countably closed follows immediately from the definition. Denote by $\overline{X}$ the union of the approximations found in G. It is not hard to see by genericity that $\overline{X}$ is a weakly uniformly $\aleph_0$-branching branchwise-real tree, and that the union of the partial rank functions yields a total rank function $\overline{X} \to \omega$.

To see that $\overline{X}$ is rigid, take $\dot{f} : \overline{X} \to \overline{X}$ a non-trivial automorphism. By a countable-closure argument, we can find a condition $\langle X, \rho, S \rangle$ forcing (i) that $\dot{f}$ is a monotonic function, (ii) that $X \subseteq \mathrm{dom}(\dot{f})$ and that $\dot{f}$ has a decided value on every element of X and (iii) that there is $x \in X$ such that $\dot{f}(x)$ is incomparable with x. Moreover we can assume, using another countable-closure argument, that there is $z \geqslant x$ in X such that both z and $\dot{f}(z)$ are branching, and furthermore that there is $w > z$ such that $[z, w]^X$ is order-dense as subset of $[z, w]^{\overline{X}}$. Since $[z, w]^X$ is countable, it has continuum-many 'holes' to fill. Any element which fills a hole has its image under $\dot{f}$ fixed, given then density of $[z, w]^X$ in $[z, w]^{\overline{X}}$. Finally, it is dense (in the $\mathbb{P}$-forcing sense) that some hole in $[z, w]^X$ is filled with a branching node while the corresponding hole in $[\dot{f}(z), \dot{f}(w)]^X$ is filled with a non-branching node which is promised to remain non-branching. Since G is generic, this means that $\dot{f}$ cannot be an automorphism of $\overline{X}$, which is a contradiction. ↯ □

2.5 A uniformly branching, rigid branchwise-real tree

In this section, I strengthen the result of the previous, by showing that for every κ with $2 \leqslant \kappa \leqslant \mathfrak{c}^+$ there is a rigid branchwise-real tree in which every point is branching of the same degree κ.

Definition 2.20. A tree order is *uniformly κ-branching* if every node has degree κ.

Moreover, we can ask that such a tree be either continuously gradable (Theorem 2.21) or not (Theorem 2.29).

Theorem 2.21. *Let $2 \leqslant \kappa \leqslant \mathfrak{c}^+$. There is a uniformly κ-branching, rigid branchwise-real tree which is continuously gradable.*

As in the proof of Theorem 2.18, our uniformly branching trees will be grown iteratively upwards. But notice that, since we require every point to have the same degree κ, at successor stages there are no choices to make: we must add $(\kappa - 1)$-many new spines above every point added at the previous stage. Were we to stop this process after ω-many steps, we would end up with a 'minimal' uniformly branching tree, which is not only non-rigid, but moreover homogeneous, as the following result shows.

Definition 2.22. Let κ be a cardinal. The *minimal uniformly κ-branching branchwise-real tree*, M_κ, is the branchwise-real tree built in ω-many steps, starting with the root, and at successors stages adding $(\kappa - 1)$-many new open spines above every point added at the previous stage.

Remark 2.23. The tree $M_\mathfrak{c}$ is the (underlying partial order of) the nowhere separable metric space constructed by Urysohn in [Ury27].

Proposition 2.24.

(1) M_κ is minimal: for every uniformly κ-branching branchwise-real tree X there is a monotonic embedding $f\colon M_\kappa \to X$. Moreover, we can assume that f is continuous: for any $x < y$ in M_κ the restriction $f{\restriction}[x, y]\colon [x, y] \to [f(x), f(y)]$ is surjective.

(2) M_κ is homogeneous: if $x, y \in M_\kappa$ are not the root, there is an order-automorphism $f\colon M_\kappa \to M_\kappa$ such that $f(x) = y$.

Proof. First note the construction of M_κ yields a rank function $\rho\colon M_\kappa \to \omega$.

(1) We can build a monotonic embedding $M_\kappa \to X$ using a Zorn's Lemma style argument, essentially following the proof of Lemma 2.9.

(2) First take any $z \in M_\kappa$ such that $n = \rho(z) > 1$. We define an automorphism $r_z\colon M_\kappa \to M_\kappa$ which steps down z's rank: $\rho(r_z(z)) = n - 1$. Let w be the greatest element of $\downarrow(z)$ of rank $n - 1$. Then z lies on a spine added above w at stage n. Let C be the component above w which contains z (which component

is of rank n). Let D be the component above w which contains the rest of the rank-$(n-1)$ spine on which w lies (which is of rank $n-1$). See Figure 2.7 for a picture of the situation. It is easy to define an isomorphism $s\colon C \to D$ such that the rank-n spine on which z lies is mapped to the rest of the rank-$(n-1)$ spine above w. Letting r_z be the automorphism which swaps C and D via s, we see that $\rho(r_z(z)) = n-1$.

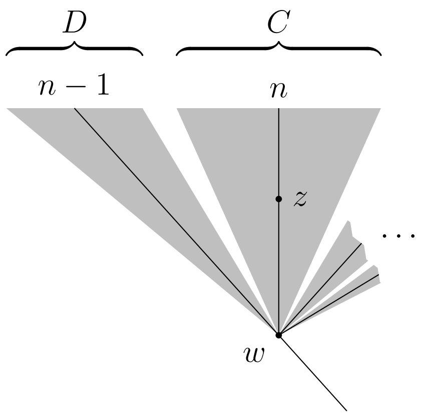

Figure 2.7: The situation when reducing z's rank. The point z lies on a rank-n spine added above w, in the component C above w. The rank-$(n-1)$ spine on which w lies continues into the component D above w. The automorphism r_z swaps the components C and D.

Now take $x, y \in M_\kappa$ non-root. By repeating the above procedure, we can assume that both have rank 1. Moreover, by performing another 'twist', sending the component above the root containing y to that containing x, we may assume that x and y are comparable, lying on the same spine; say $x \leqslant y$. Finally, it is not hard to define an automorphism of M_κ whose restriction to the rank-1 spine containing the points x and y looks like an automorphism of $(0, \infty)$, and which maps x to y. □

Turning back to the construction of rigid trees, in the proof of Theorem 2.21, as well as Theorem 2.29 below, the notion of a 'colouring' of the positive real numbers $(0, \infty)$ plays an auxiliary role. A colouring of the positive real numbers is simply a function with domain $(0, \infty)$; we think of elements of the range as 'colours'.

Definition 2.25. Let X and S be sets. An *S-colouring* of X is a function $X \to S$.

In the proof of Theorem 2.18, the first step was essentially to construct a family of black-white colourings of $(0, \infty)$ with suitable properties. To construct a uniformly branching, rigid branchwise-real tree, we extend this idea. This time we look for a family of 2-colourings (and later ω-colourings) which is 'sufficiently generic'. This means that for any pair of distinct colourings, any order-automorphism of $(0, \infty)$ and any pair of colours, we can densely often find a point coloured with the first colour under the first colouring whose image under the automorphism is coloured with the second colour under the second colouring. This is the following result.

Lemma 2.26 (Generic Colouring Lemma)**.** *Let $\lambda \leqslant \mathfrak{c}$ be a cardinal. There is a family $\mathcal{A}$ with size $\mathfrak{c}^+$ of λ-colourings $(0, \infty) \to \lambda$ of the positive real numbers such that the following holds. For any $c, d \in \mathcal{A}$ distinct, for any order-automorphism $p\colon (0, \infty) \to (0, \infty)$ and for any $\alpha, \beta \in \lambda$ there is a dense set of $x \in (0, \infty)$ such that $c(x) = \alpha$ and $d(p(x)) = \beta$.*

To prove this we need the following basic result concerning automorphisms of the positive real numbers.

Lemma 2.27. *There are exactly $\mathfrak{c}$-many order-automorphisms $(0, \infty) \to (0, \infty)$.*

Proof. Every real number is the limit of rationals below it, so any order-automorphism $(0, \infty) \to (0, \infty)$ is determined by its values on $(0, \infty) \cap \mathbb{Q}$, of which there are $\mathfrak{c}$-many possibilities. □

Proof of Lemma 2.26. By Zorn's Lemma, we can take a maximal family $\mathcal{A}$ of λ-colourings, satisfying both the property in the statement, as well as that for every $c \in \mathcal{A}$:

$$\text{for every } \alpha \in \lambda \text{ and for every } b > a > 0 \text{ there are continuum-many points } x \in (a, b) \text{ such that } c(x) = \alpha \tag{P}$$

Let us first see that $\mathcal{A}$ is non-empty. For convenience, I will construct a colouring of $(0, 1)$ satisfying (P); this can then easily be adapted to a colouring of $(0, \infty)$. It suffices to take $\lambda = \mathfrak{c}$. Consider the elements of $(0, 1)$ as binary ω-sequences representing binary expansions (where, to avoid ambiguity, we make some canonical choice for the binary representation of $x \in (0, 1)$, in the case that two such representations are possible). Define an equivalence relation on these, where two sequences are related if one is a tail of the other. Each equivalence class is countable, hence there are continuum-many. Moreover, each equivalence class is dense in $(0, 1)$. Group the classes into continuum-many batches of size $\mathfrak{c}$, and colour all of each batch with a different element of $\mathfrak{c}$. The resulting colouring satisfies (P).

Now suppose for a contradiction that $|\mathcal{A}| < \mathfrak{c}^+$. We will extend $\mathcal{A}$ by diagonalising against the previous colourings, and against every automorphism. Let $\operatorname{Aut}(0, \infty)$ be the set of order-automorphisms $(0,\infty) \to (0,\infty)$. By Lemma 2.27, the set $\mathrm{Aut}(0,\infty) \times \mathcal{A} \times \lambda \times \lambda$ has size $\mathfrak{c}$. Enumerate:

$$\mathrm{Aut}(0,\infty) \times \mathcal{A} \times \lambda \times \lambda = \{(p_\theta, c_\theta, \alpha_\theta, \beta_\theta) \mid \theta < \mathfrak{c}\}$$

We build a new colouring $d\colon (0,\infty) \to \lambda$ recursively in stages d_θ for $\theta < \mathfrak{c}$, each of which is a partial colouring with domain of size less than $\mathfrak{c}$. Assume we have constructed d_μ for $\mu < \theta$. First let $d'_\theta := \bigcup_{\mu<\theta} d_\mu$. Now, since c_θ satisfies (P), and since d'_θ has domain of size less than $\mathfrak{c}$, we can find a countable dense set of $X \subseteq (0,\infty)$ such that for every $x \in X$ we have $c_\theta(x) = \alpha_\theta$ while $d'_\theta(p_\theta(x))$ is undefined. Extend d'_θ to d_θ by letting $d_\theta(p_\theta(x)) := \beta_\theta$, for every $x \in X$. Finally, let d be the union $\bigcup_{\theta<\mathfrak{c}} d_\theta$, filling in any points which remain uncoloured arbitrarily. But now, $\mathcal{A} \cup \{d\}$ is a larger family, contradicting the maximality of $\mathcal{A}$. ↯ □

With this lemma established, we can now construct our uniformly κ-branching, rigid branchwise-real tree which is continuously gradable.

Proof of Theorem 2.21. Let $\mathcal{A}$ be a $\mathfrak{c}^+$-sized family of 2-colourings of $(0,\infty)$ as per the Generic Colouring Lemma 2.26. This time, we construct X in stages X_α for $\alpha < \omega^2$. When we add a new spine we lay a new colouring from $\mathcal{A}$ along it, and thus we also build colourings $c_\alpha\colon X_\alpha \to 2$. At limit stages, we use these colourings to decide which branches to extend. We also keep track of the rank ρ of elements of X, so that the points added at stage α get rank α. As before, we need to make sure we never run out of colourings, so we can partition $\mathcal{A}$ into ω^2-many batches of size $\mathfrak{c}^+$.

Start with X_0 the singleton, coloured 0. Assume that X_α and c_α are constructed, where α is 0 or a limit. To make $X_{\alpha+1}$, above all of the points of rank α, we add κ-many new open spines. For each spine added, we pick a different, unused colouring $c \in \mathcal{A}$, and use it to define $c_{\alpha+1}$ on that spine. Now assume that X_α and c_α are constructed, where α is a successor. To make $X_{\alpha+1}$, above each point of rank α, we add $(\kappa - 1)$-many new open spines, again colouring them with new colourings from $\mathcal{A}$.

So take $\alpha < \omega^2$ a limit, and assume we have constructed X_β and c_β for $\beta < \alpha$. First let $X'_\alpha := \bigcup_{\beta<\alpha} X_\alpha$ and $c'_\alpha := \bigcup_{\beta<\alpha} c_\alpha$. Then X'_α admits a number of new branches which appear in the limit: in other words, branches through X'_α containing elements of rank unbounded in α. It is these branches which we will decide to extend or not. Let B be any such branch. Then for every $\beta < \alpha$, there is a maximum element $x_\beta \in B$ of rank β. Consider the sequence $(c'_\alpha(x_0), c'_\alpha(x_1), \ldots)$ of the colours of such points. We will extend B if and only if this sequence has a tail consisting of 1's. To extend B, add a new point to X'_α lying directly above B, and colour it 0. Once we have carried out this procedure for every new branch through X'_α, we arrive at our new pair $\langle X_\alpha, c_\alpha \rangle$.

Finally, let $X := \bigcup_{\alpha<\omega^2} X_\alpha$ and $c := \bigcup_{\alpha<\omega^2} c_\alpha$. Note that at each stage we made sure that every point is branching of degree κ, so X is uniformly κ-branching. Moreover, $\rho\colon X \to \omega^2$ is a bounded rank function, so X is continuously gradable by Lemma 2.12.

Let us see that X is rigid. Suppose for a contradiction that $f\colon X \to X$ is a non-trivial automorphism. Then there is $x_0 \in X$ such that $f(x_0) \neq x_0$. When $\kappa \geqslant 3$, the proof is simpler, so let's deal with that case first. See Figure 2.8 for a representation of the situation. Consider the components above above x_0 and above $f(x_0)$. By Lemma 2.16, f factors through a bijection of the components above x_0 to those above $f(x_0)$. Moreover, all but at most one component above x_0 is of rank $\rho(x_0)+1$, and similarly for $f(x_0)$. Since $\deg(x_0) = \deg(f(x_0)) = \kappa \geqslant 3$ there is a component C above x_0 of rank $\rho(x_0)+1$ which maps to a component $f(C)$ above $f(x_0)$ of rank $\rho(f(x_0))+1$. Furthermore, using Lemma 2.16 again, there is an coinitial interval $I \subseteq C$ of constant rank $\rho(x_0)+1$ which maps to a coinitial interval $f(I) \subseteq f(C)$ of constant rank $\rho(f(x_0))+1$. Then I is an initial segment of a spine U added at stage $\rho(x_0)+1$, while $f(I)$ is an initial segment of a spine V added at stage $\rho(f(x_0))+1$.

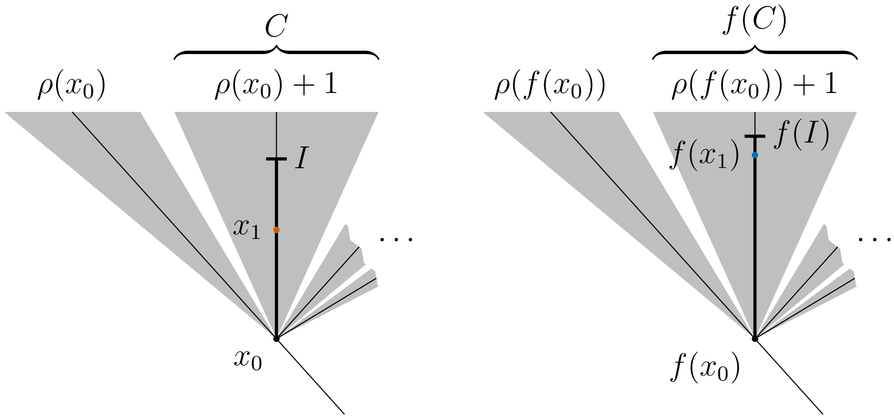

Figure 2.8: Picking x_1 when $\kappa \geqslant 3$. The grey regions represent connected components, while thin lines represent spines. The ordinal above a spine indicates its rank.

Now, consider the colouring c along I and along $f(I)$. By construction, these come from different colourings from $\mathcal{A}$. The order-isomorphism $f{\restriction}I\colon I \to f(I)$ induces an isomorphism $U \to V$, which in turn induces an automorphism of $(0,\infty)$. Hence by the key property of the family $\mathcal{A}$ guaranteed by the Generic Colouring Lemma 2.26, there is $x_1 \in I$ coloured $c(x_1) = 1$ such that its image under f is coloured $c(f(x_1)) = 0$.

Iterate this process to produce a sequence $x_0 < x_1 < \cdots$ such that $c(x_n) = 1$ and $c(f(x_n)) = 0$ for every $n > 0$. Then $\{x_0, x_1, \ldots\}$ determines a branch B through $X'_{\rho(x_0)+\omega}$, while $\{f(x_0), f(x_1), \ldots\}$ determines a branch $f(B)$ through $X'_{\rho(f(x_0))+\omega}$. Moreover, when deciding whether to extend B at stage $\rho(x_0) + \omega$, we use a sequence whose tail is $(x_0, x_1, \ldots)$, and respectively for $f(B)$. But then the former branch gets extended, while the latter does not, contradicting that f is an automorphism. ↯

Let us now turn to the case $\kappa = 2$. The issue here is that we can no longer guarantee that an initial segment of a new spine added above x_0 gets mapped to an initial segment of a new spine added above $f(x_0)$. Here's how we proceed. If the component C above x_0 of rank $\rho(x_0) + 1$ is mapped to the component above $f(x_0)$ of rank $\rho(f(x_0)) + 1$, proceed as before. Otherwise, the rank of $f(C)$ must be $\rho(f(x_0))$. See Figure 2.9 for a representation of the situation. By Lemma 2.16 there is a coinitial interval $I \subseteq C$ of constant rank whose image $f(I) \subseteq f(C)$ has constant rank. So I is an initial segment of a spine added above x_0 at stage $\rho(x_0) + 1$ which maps onto a segment of a spine which already exists at stage $\rho(f(x_0))$. Now, using the key property of the colouring $\mathcal{A}$, pick any element x'_0 in the interior of I with colour $c(x'_0) = 1$ whose image has colour $c(f(x'_0)) = 0$. Then we must have that the component C' above x'_0 of rank $\rho(x'_0) + 1$ is mapped to the component above $f(x'_0)$ of rank $\rho(f(x'_0)) + 1$, and so we can proceed as before to obtain $x_1 > x'_0$ coloured $c(x_1) = 1$ whose image is coloured $c(f(x_1)) = 0$. Iterating, we again build a sequence $x_0 < x'_1 < x_1 < \cdots$ through X, all of whose elements (except possibly x_0) are coloured 1 but with image coloured 0. When deciding whether to extend the corresponding branch through $X'_{\rho(x_0)+\omega}$, we use the colours of a subsequence of $(x_0, x'_1, x_1, \ldots)$. In the end, we find that this sequence has a limit in X, while its image does not. ↯ □

With this ZFC result established, let us examine once more the connection with a forcing notion. We again work with the $\kappa = \aleph_0$ case. We will define forcing conditions similarly to the discussion at the end of Section 2.4. This time however, instead of making promises that some points won't become branching, we promise that certain branches won't get ever get extended. There are a number of different ways of achieving this. (1) The most direct way is to include in each condition a set of branches through the sub-branchwise-real tree, which are promised never to be extended. A technical issue which arises is that a branch through a smaller sub-branchwise-real tree need not be a branch through a larger one, since it need not be maximal as a chain any more. To make sure that the set of promises remains valid, we require that a stronger condition include a promise set which contains the unique extension of each promised branch in the weaker tree to the larger one. (2) A more elegant method is to keep track instead

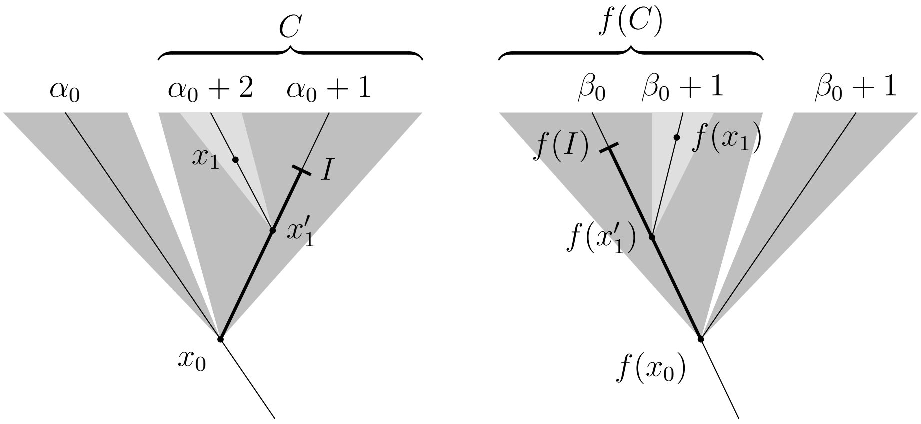

Figure 2.9: Picking x_1 when $\kappa = 2$ and $\rho(f(C)) = \rho(f(x_0))$. The dark and light grey regions represent connected components, while thin lines represent spines. The ordinal above a spine indicates its rank, where $\alpha_0 = \rho(x_0)$ and $\beta_0 = \rho(f(x_0))$.

of a monotonic function from the approximations into $\mathbb{R}$. This way the branches which are not to be extended are exactly those on which the monotonic function is unbounded. (3) Yet a third way is make use of the universal continuously-gradable branchwise-real tree U_κ defined in Remark 2.14. Any maximal antichain A through U_κ defines a uniformly κ-branching branchwise-real tree: take $\{x \in U_\kappa \mid \exists a \in A \colon x < a\}$. We can then consider simply the set of all antichains in U_κ under extension as a forcing poset. For the following result, I remain agnostic as to the exact method used.

Theorem 2.28. *A generic tree $\overline{X}$ for the forcing just described is a uniformly κ-branching, rigid branchwise-real tree.*

Proof sketch. The hard part is again showing that $\overline{X}$ is rigid. Take $\dot{f} \colon \overline{X} \to \overline{X}$ a non-trivial automorphism. By a countable-closure argument, we can find a condition p_0 forcing that $\dot{f}$ is a non-trivial monotonic function whose domain contains the approximation X_0 at p_0 with $\dot{f}$ decided on every element of X_0, and with an element $x_0 \in X_0$ such that $\dot{f}(x_0)$ is incomparable with x_0. By another countable-closure argument, it is dense to find a condition $p_1 < p_0$ such that the domain of $\dot{f}$ contains the approximation X_1 at p_1 and such that there is $x_1 \in X_1$ not less than any element in X_0 with $x_0 < x_1$. Continuing in this way, we build a chain of forcing conditions $p_0 > p_1 > \cdots$ together with $x_0 < x_1 < \cdots$ in $\bigcup_{n\in\omega} X_n$. Let p_ω be the limit of the p_n's. Then we can strengthen p_ω so that the branch determined by $\{x_0, x_1, \ldots\}$ gets extended, while we promise never to extend the branch determined by $\{\dot{f}(x_0), \dot{f}(x_1), \ldots\}$. (In the case where we keep track of a monotonic function into $\mathbb{R}$, we would need to make a tweak to ensure that one sequence remains bounded while the other becomes unbounded.) Therefore, we can densely find a contradiction to the fact that $\dot{f}$ is monotonic. $\lightning$ □

Turning back to our ZFC constructions, the tree constructed in Theorem 2.21 is continuously gradable. By suitably modifying the method, we can construct one which is not continuously gradable.

Theorem 2.29. *Let $2 \leqslant \kappa \leqslant \mathfrak{c}^+$. There is a uniformly κ-branching, rigid branchwise-real tree which is not continuously gradable.*

This time, the construction must proceed for ω_1-many steps, because Lemma 2.12 shows that any construction terminating after countably many steps would yield a continuously gradable tree. This then introduces a complication, as we need to ensure that there are no branches through the final tree which contain points of every rank below ω_1. Indeed, such a branch would contain an ω_1-sequence, violating condition (BR1) that every branch is isomorphic to a real interval. Now, as mentioned in Remark 2.14, not every tree constructed in ω_1-many steps is non-continuously-gradable, so we also need to guarantee this. This is done by realising a known non-$\mathbb{R}$-gradable well-stratified tree inside the final branchwise-real tree, following techniques used in Chapter 1. Lastly, these two competing requirements need to be balanced with the need for the final tree to be rigid. These desiderata are achieved following the method used in the proof of Theorem 2.21 above, this time utilising ω-gradings instead of 2-gradings. At limit steps, the idea is essentially to extend a branch if and only if in the sequence of colours $(c'_\alpha(x_0), c'_\alpha(x_1), \ldots)$ of the maximal points of each rank, every element of ω appears only finitely often.

First, let us meet the non-$\mathbb{R}$-gradable well-stratified tree without uncountable branches which will be found inside the final branchwise-real tree.

Lemma 2.30. *Let T be the tree of functions r into ω, whose domain is a countable ordinal, such that when restricted to the set of successor ordinals r is injective into $\omega \setminus \{0\}$, and elsewhere r is identically 0. We consider T as a partial order under function extension. Then T is a well-stratified tree without an $\mathbb{R}$-grading.*

Proof. The tree T is isomorphic to the tree $\mathrm{Pad}(\mathrm{In}_\omega)$ defined in Section 1.3. There, using a proof due to Baumgartner, Gavin and Laver [Bau70], it is shown that this tree is well-stratified, has no $\mathbb{R}$-grading, and moreover that every branch is countable. □

With this to hand, we can now prove the theorem.

Proof of Theorem 2.29. Let $\mathcal{A}$ be a $\mathfrak{c}^+$-sized family of $(\omega \setminus \{0\})$-colourings of $(0, \infty)$ as per the Generic Colouring Lemma 2.26. We construct X in stages X_α for $\alpha < \omega_1$, colouring each spine appearing at a stage which is the successor of a successor using a different colouring from $\mathcal{A}$. As in Theorem 2.21, these colourings are used to determine which branches to extend at limit stages. This will ensure that the resulting tree X is rigid, that every branch is isomorphic to a real interval and that there is no continuous grading.

Proceed as in Theorem 2.21. First partition $\mathcal{A}$ into ω_1-many batches of size $\mathfrak{c}^+$, so that we never run out. We build each X_α and colouring c_α, together with the rank function ρ, recursively. Start with X_0 the singleton coloured 0. As before, at successor stages add $(\kappa - 1)$-many new open spines above the points added at the previous stage. If we are at the successor of 0 or a limit, colour the new spines all with 0, otherwise colour each using new colourings from $\mathcal{A}$. Now assume that $\alpha < \omega_1$ is a limit. First let $X'_\alpha := \bigcup_{\beta<\alpha} X_\beta$ and $c'_\alpha := \bigcup_{\beta<\alpha} c_\beta$. We need to decide which new branches through X'_α to extend. As before, a branch B is determined by its sequence $(x_0, x_1, \ldots)$ of maximum points of each rank $\beta < \alpha$. Extend the branch B if and only if no non-zero colour appears infinitely often in the sequence $(c'_\alpha(x_0), c'_\alpha(x_1), \ldots)$. To extend a branch, add a new point above it coloured 0. Let X_α and c_α be result of performing this operation on every branch in X'_α.

Finally, let $X := \bigcup_{\alpha<\omega_1} X_\alpha$ and $c := \bigcup_{\alpha<\omega_1} c_\alpha$. Note that at each stage we made sure that every point is branching of degree κ, so X is uniformly κ-branching. To see that X is branchwise-real, take any branch B through X. Then B is determined by its sequence $(x_0, x_1, \ldots)$ of maximum points of each rank $\beta < \omega_1$. Consider the sequence $(c_\alpha(x_0), c_\alpha(x_1), \ldots)$. Every colour indexed by the successor of a successor is non-zero, and every non-zero colour appears at most countably many times. Hence the sequence must have countable length. Therefore, the branch B cannot contain a subset isomorphic to ω_1, so it must be isomorphic to a real interval, as required.

To show that X is rigid, we follow the method in Theorem 2.21. If $f\colon X \to X$ is a non-trivial automorphism, we can pick x_0 such that $f(x_0) \neq x_0$. Note that if we consider the sequence of points below x_0, respectively $f(x_0)$, maximal in each rank, then no non-zero colour appears infinitely often. So to determine whether a branch containing x_0 gets extended at a limit, it suffices to consider the sequence of colours of points maximal in each rank lying *above* x_0, and likewise for $f(x_0)$. As in Theorem 2.21 we can construct a sequence $x_0 < x_1 < \cdots$ recursively, this time ensuring that $c(x_n) = n$ while $c(f(x_n)) = 1$, for each $n > 2$. Note that while any of $x_0, x_1, f(x_0), f(x_1)$ may be coloured 0, none of the rest of the points or their images can be, since they all have ranks which are the successor of a successor. Then the sequence $(x_0, x_1, \ldots)$ has a limit in X, while its image does not.

Finally, to see that X has no $\mathbb{R}$-grading, we realise the tree T given by Lemma 2.30 as a subtree. For this, we recursively construct an embedding $\iota\colon T \to X$ such that for every $r \in T$:

- $\rho(\iota(r)) = \mathrm{rank}(r)$, and
- $c(\iota(r)) = r(\alpha)$ when $\mathrm{rank}(r) = \alpha + 1$.

Start by sending the root $\varnothing$ to the root.

Take $r \in T$ and assume that we have defined $\iota(r)$. There are two cases. If $\mathrm{rank}(r)$ is 0 or a limit, then r has exactly one immediate successor s. Moreover $\rho(\iota(r))$ is 0 or a limit, so by the way we constructed X there are κ-many spines with base $\iota(r)$ added at stage $\mathrm{rank}(r) + 1$, all of which are coloured 0. Let $\iota(s)$ be any point on any one of these spines. The other case is when $\mathrm{rank}(r)$ is a successor. Then r has κ-many successors. Moreover, there are κ-many spines with base $\iota(r)$ added at stage $\mathrm{rank}(r)+1$, each of which are coloured according to a colouring from $\mathcal{A}$. In particular, every non-zero colour appears on every spine. Place each successor s of r on a different spine, in such a way that $c(\iota(s)) = s(\mathrm{rank}(r))$.

Now take $r \in T$ with $\mathrm{rank}(r)$ a limit, and assume that we have defined $\iota(s)$ for all $s < r$. Consider the sequence $(\iota(r{\restriction}0), \iota(r{\restriction}1), \ldots)$. This determines a branch through $X'_{\mathrm{rank}(r)}$. Moreover, when determining whether to extend this branch, the sequence of colours used is:

$$(c(\iota(r{\restriction}0)), c(\iota(r{\restriction}1)), \ldots) = (0, r(0), r(1), \ldots, 0, r(\omega), r(\omega+1), \ldots)$$

By definition of T, no non-zero colour appears more than once in this sequence. Therefore, the branch is extended, meaning that the sequence $(\iota(r{\restriction}0), \iota(r{\restriction}1), \ldots)$ has a supremum, which we can set as the image of r.

Putting it all together, we obtain an embedding $\iota\colon T \to X$. Any continuous grading of the latter would restrict to an $\mathbb{R}$-grading of the former, which by Lemma 2.30 is impossible. □

2.6 Open questions

In this final section I present a number of questions left open by the preceding investigation.

The strongest result in this chapter is that there exists both a continuously gradable and a non-continuously-gradable uniformly κ-branching, rigid branchwise-real tree for $2 \leqslant \kappa \leqslant \mathfrak{c}^+$. This limit of $\mathfrak{c}^+$ comes from the Generic Colouring Lemma 2.26. It is natural to wonder if we could get past this limit using some alternative strategy.

Question 2.31. *Does there exist a uniformly κ-branching, rigid branchwise-real tree for $\kappa > \mathfrak{c}^+$?*

Now, it may be that such a question is beyond ZFC, and that the best we can hope for is a consistency result. We might try expanding the Generic Colouring Lemma by forcing both a $2^{\mathfrak{c}}$-sized family of colourings and that $2^{\mathfrak{c}} > \mathfrak{c}^+$. Or alternatively we could try to force our rigid tree directly, for example by using a forcing notion similar to that presented in Section 2.5.

Taking a different tack, throughout the present chapter we have been aiming to eliminate all order-automorphisms. What if wanted more fine-grained control over the automorphism group? It is not hard to modify the methods presented here to yield a branchwise-real tree with automorphism group (i) the symmetries of κ as a set, (ii) $\mathbb{Z}$, (iii) the order-automorphism group of the positive reals or (iv) various combinations of these.

Question 2.32. *What is the class of automorphism groups of (uniformly κ-branching) branchwise-real trees?*

Finally, we have seen four examples of uniformly κ-branching branchwise-real trees: the minimal M_κ, the universal U_κ, the rigid continuously gradable one and the rigid non-continuously-gradable one. None of these trees are isomorphic. Moreover, each comes with a rank function, the supremum of whose values is, respectively, ω, ω_1, ω^2 and ω_1. Clearly, any branchwise-real tree may admit many different rank functions, but is there some canonical way in which we might stratify uniformly κ-branching branchwise-real trees in terms of their 'complexity'? One direction of investigation might be to first investigate how these tree embed into one another. We already know that M_κ is a minimum and U_κ a maximum, as least for the continuously gradable ones.

Question 2.33. *What is the nature of the class of uniformly κ-branching (continuously gradable) branchwise-real trees under embeddability?*

Alternatively, we might consider the following isomorphism-invariant of branchwise-wise real trees X: what is the minimum supremum value of a rank function on X? For

M_κ, this is ω, and for any non-continuously gradable branchwise-real tree, this must be ω_1 (by Lemma 2.12). It is not immediate that the uniformly κ-branching, rigid, continuously gradable branchwise-real tree constructed in Section 2.5 has minimum supremum rank ω^2; however it is not hard to see that its minimum cannot be ω. This leads to the following questions.

Question 2.34. *Which ordinals arise as the minimum supremum value of a rank function on a uniformly κ-branching branchwise-real tree?*

Chapter 3

Bisimulations of potentialist systems

3.1 Introduction

The two main objects of study in this chapter are potentialist systems and bisimulations. A potentialist system is a collection of first-order structures and embeddings between them in which one can interpret first-order modal logic (for definitions, see Section 3.2). Potentialist systems are closely related to first-order Kripke models. The former differ from the latter in the following three ways. (1) All instances of the accessibility relation between structures are required to be embeddings; in other words, the accessibility relation is required to be inflationary on the domains. (2) We allow for the existence of multiple embeddings between two structures. (3) We allow class-many structures.

Potentialist systems provide a way of investigating a mathematical structure within the context of other related structures, making precise the notions of necessity and possibility. A formal treatment of these systems is given in [HL19]. The study of potentialist systems is called 'modal model theory'. This name was introduced in [HW20], which conducts a thorough study of the system $\mathrm{Mod}(T)$ of models of a first-order theory T. This system, and more general questions relating to potentialist systems, were also studied in [SS20] (using different terminology). Potentialist systems of set-theoretic forcing extensions are studied in [Ham03; HL08; HL13; HL19; HLL15]. Further systems of models of arithmetic and set theory are studied in [HW18; Ham18; HW21].

A bisimulation between two potentialist systems is a system of finite partial isomorphisms between structures in one and structures in the other, such that any increase in the size of the isomorphism or travel along the accessibility relation on one side can be mirrored on the other. A bisimulation is bitotal if every element of every model in both systems takes part in the relation. This notion of bisimulation simultaneously generalises the notion of a bisimulation between propositional Kripke models and the notion of a back-and-forth system for first-order models. One of the key properties which bisimulations satisfy is invariance: if two systems are bisimilar, they satisfy the same modal sentences, even allowing for infinite disjunctions and conjunctions.

An important class of examples of potentialist systems are those of the form $\mathrm{Mod}(T)$, consisting of every model of some first-order theory T. There are two ways of defining $\mathrm{Mod}(T)$: by including all embeddings between models, or just the substructure inclusions. As discussed in [HW20], the latter aligns well with a potentialist philosophy on the theory T, in which we view the 'universe' as unfolding in stages, with individuals becoming actual and persisting identically through subsequent stages. The former on the other hand typically enjoys better algebraic properties, like directedness, amalgamation or convergence (see [HW20, Section 5] for definitions). In Theorem 29 of that paper, Hamkins and Wołoszyn show that the two resulting systems satisfy the same modal sentences. I go one step further and show that the two systems are in fact bitotally bisimilar in a particularly strong way (Theorem 3.28). This provides a deeper explanation for why the two are equivalent on modal sentences, and immediately extends the result to equivalence with respect to an infinitary language. Moreover, the bisimulation should allow us to transfer results from the more mathematically amenable system based on embeddings, to the more philosophically relevant one based on inclusions.

Bisimulations have a natural interpretation in terms of infinite games. I elaborate on this connection, defining the notion of a *modal Ehrenfeucht-Fraïssé game*. This furnishes us with the converse of invariance: if two systems agree on all infinitary modal sentences then they are bisimilar.

For the remainder of the chapter, I consider under what circumstances a class-sized potentialist system is bitotally bisimilar with a set-sized one. First, I restrict attention to the systems $\mathrm{Mod}(T)$ for some first-order theory T. I show that if T has infinite models, and all its models of size κ are $\aleph_0$-saturated, then $\mathrm{Mod}(T)$ is bitotally bisimilar to a system consisting of set-many models. In the last part I take a different tack, showing that if a potentialist system consists of models of an $\aleph_0$-categorical theory which admits quantifier elimination then it is bitotally bisimilar to a particularly simple system containing set-many models. This result makes no assumption on the modal structure of the system, and as such it might be expected that it wouldn't admit much strengthening. As evidence for this, I provide a partial converse, and give an example where weakening the assumption 'admits quantifier elimination' leads to a system which is not bisimilar with a set-sized one.

3.2 Potentialist systems

A propositional Kripke model consists of a set of 'worlds', a relation on this set, and an assignment to each world of the truth values of each propositional variable [BRV01; CZ97]. These models are used to provide a semantics for propositional modal logic. They generalise to first-order Kripke models, in which each world is now an entire first-order structure, and we specify how the elements of connected worlds are related. See [HC96; FM98; GSS09; Ben10] for more on first-order Kripke models.

Potentialist systems specialise first-order Kripke models, in that we require that all relations between worlds be embeddings. However, they also generalise, in that we allow multiple relations between any two worlds, and that the collection of worlds be class-sized. The terminology 'potentialist system' was first used in [HL19].

Definition 3.1. A *potentialist system* is a collection of structures in a given vocabulary together with a collection of embeddings between them, containing the identity embeddings and all compositions of the embeddings.

I will refer to the elements of a potentialist system as 'worlds', 'structures' or 'models' (the latter being reserved for the context in which we are considering models of a certain first-order theory). Note that a potentialist system is exactly a category of first-order structures with embeddings. A potentialist system is an instance of the more general notion of a 'Kripke category', in which one considers the modal logic of an arbitrary concrete category. Kripke categories are defined and investigated in upcoming work by Wojciech Aleksander Wołoszyn [Woł22].

Notice that we allow multiple embeddings between worlds. Sometimes, it is natural to consider potentialist systems in which embeddings are always unique, so the abstracted relationship between the worlds is a preorder.

Definition 3.2. A potentialist system is *thin* if from any structure in it to any other there is at most one embedding.

This corresponds with the notion of 'thinness' for categories. In Theorem 3.28 below I will show that restricting to thin potentialist systems does not lose generality, in the sense that every potentialist system is bisimilar with a thin one.

Definition 3.3. A *pointed system* is a triple $(\mathcal{M}, M, \overline{a})$, where $\mathcal{M}$ is a potentialist system, $M \in \mathcal{M}$ and $\overline{a}$ is a finite tuple of elements of M. The *parameter size* of the pointed system is the size of $\overline{a}$. When the parameter size is 0, write $(\mathcal{M}, M)$.

Let us now move on to consider the interpretation of formulas in potentialist systems. We are interested in several languages. We start with some first-order, finitary, non-modal language L (specified by its signature of non-logical symbols). The first language to meet is the standard first-order modal language based on L.

Definition 3.4. The language $L^{\Diamond}$ is formed by adding $\Diamond$ and $\Box$ as operators. In other words, $L^{\Diamond}$ is closed under boolean combinations, quantification and the follow two new rules: whenever ϕ is a formula, so are $\Diamond\phi$ and $\Box\phi$.

Now, when we come to consider the relationship between bisimulations and sentence satisfaction, we will need to talk about infinitary languages. The definition generalises that of infinitary first-order (non-modal) languages [Hod93] (see also [SS20]).

Definition 3.5. For κ a regular, infinite cardinal, let L_κ be the first-order language in which we allow conjunction and disjunction of size less than κ, with the proviso that all formulas contain only finitely many free variables. Let L_∞ be the (class-sized) union of each L_κ. The languages $L_\kappa^{\Diamond}$ and $L_\infty^{\Diamond}$ are defined similarly, this time including $\Diamond$ and $\Box$ as modal operators.

Remark 3.6. In non-modal first-order infinitary languages, one can also consider allowing infinitely long quantifier blocks. With this in mind the language L_κ is elsewhere commonly denoted by $L_{\kappa,\omega}$, where the 'ω' subscript makes explicit the fact that quantifier blocks are all finite; likewise L_∞ is written $L_{\infty,\omega}$. Similarly, in the modal case, one might consider allowing infinitely long blocks of modal operators. So, strictly speaking, we could write $L^{\Diamond}_{\kappa,\omega,\omega}$ for $L^{\Diamond}_\kappa$ and $L^{\Diamond}_{\infty,\omega,\omega}$ for $L^{\Diamond}_\infty$. However, since infinitely long quantifier or modal operator blocks are not considered here, I suppress the additional subscripts.

Finally, since formula trees are well-founded, one can assign a rank to each formula in $L^{\Diamond}_\infty$. This provides a useful complexity stratification, which will be relevant when we come to consider modal Ehrenfeucht-Fraïssé games.

Definition 3.7. The *modal-quantifier rank*, $\mathrm{mqr}(\phi)$, of a formula $\phi \in L_\infty^\Diamond$ is an ordinal defined recursively on the construction of ϕ.

$$\begin{aligned} \mathrm{mqr}(\psi) &= 0 && \text{(when } \psi \text{ is atomic)} \\ \mathrm{mqr}(\neg\phi) &= \mathrm{mqr}(\phi) \\ \mathrm{mqr}\left(\bigvee_{i\in I}\phi_i\right) = \mathrm{mqr}\left(\bigwedge_{i\in I}\phi_i\right) &= \sup_{i\in I}\mathrm{mqr}(\phi_i) \\ \mathrm{mqr}(\exists x\phi) = \mathrm{mqr}(\forall x\phi) &= \mathrm{mqr}(\phi)+1 \\ \mathrm{mqr}(\Diamond\phi) = \mathrm{mqr}(\Box\phi) &= \mathrm{mqr}(\phi)+1 \end{aligned}$$

Write $L_\infty^{\Diamond;\alpha}$ for the class of $L_\infty^\Diamond$-formulas of modal-quantifier rank less than α.

The interpretation of first-order modal formulas extends that of classical formulas by interpreting $\Diamond\phi$ as 'ϕ holds in some extension' and $\Box\phi$ as the dual: 'ϕ holds in every extension'.

Definition 3.8. Let $(\mathcal{M}, M, \overline{a})$ be a pointed system, and ϕ be an $L_\infty^\Diamond$-formula. We define what it means for ϕ to be true at this pointed system, which we write as $\mathcal{M}, M \vDash \phi[\overline{a}]$, by induction on the complexity of ϕ. The non-modal cases are as usual. We let $\mathcal{M}, M \vDash \Diamond\phi[\overline{a}]$ if and only if there is $\pi\colon M \to M'$ in $\mathcal{M}$ such that $\mathcal{M}, M' \vDash \phi[\pi(\overline{a})]$. We let $\mathcal{M}, M \vDash \Box\phi[\overline{a}]$ if and only if for every $\pi\colon M \to M'$ in $\mathcal{M}$ we have $\mathcal{M}, M' \vDash \phi[\pi(\overline{a})]$.

Note that $\Box\phi$ is equivalent to $\neg\Diamond\neg\phi$, so in our inductions on formulas we won't need to treat the $\Box\phi$ case separately.

Remark 3.9. There are some definability issues here, stemming from the fact that the collection of worlds in $\mathcal{M}$ may be a proper class. For instance, if we take $\mathcal{M}$ to be the class of all initial segments V_α of the cumulative hierarchy under subset inclusion, then $L^\Diamond$ in $\mathcal{M}$ can capture first-order truth in the ambient universe (see [HL19, Theorem 1] and [Lin13]). This issue is discussed in Section 9 of [HW20]. Following the discussion there, to make sense of the previous definition we can view it as taking place within Morse-Kelley set theory, or von Neumann-Bernays-Gödel set theory with the axiom of elementary transfinite recursion. Alternatively, we can choose to work with a sequence of Grothendieck-Zermelo universes, so that the extension of some potentialist system in one universe is a set in the next universe up. Note that ZFC-definability issues don't arise for all class-sized systems. Indeed, Sections 3.5 and 3.6 provide examples of class-sized systems which are strongly equivalent to set-sized ones.

Definition 3.10. The relations $\equiv$, $\equiv^{\diamond}$, $\equiv^{\diamond}_{\infty}$ and $\equiv^{\diamond,\alpha}_{\infty}$ between pointed systems are the elementary equivalence relations with respect to L, $L^{\diamond}$, $L^{\diamond}_{\infty}$ and $L^{\diamond,\alpha}_{\infty}$ respectively, taking the tuples in the pointed systems as parameters. For example, $(\mathcal{M}, M, \overline{a}) \equiv^{\diamond}_{\infty} (\mathcal{N}, N, \overline{b})$ means that for every $L^{\diamond}_{\infty}$-formula ϕ, we have $\mathcal{M}, M \vDash \phi[\overline{a}]$ if and only if $\mathcal{N}, N \vDash \phi[\overline{b}]$.

Let us meet an important class of potentialist systems, studied in [HW20] and [SS20]. Let T be a first-order L-theory. The potentialist system $\mathrm{Mod}(T)$ will be the system based on the collection of all models of T. There are two ways of defining this, depending on whether we take the relations between structures to be inclusions or embeddings.

Definition 3.11. The *substructure potentialist system*, $\mathrm{Mod}(T)$, based on T, consists of the class of models of T together with all substructure inclusions. That is, we include all and only the substructure inclusion maps $\iota\colon M \subseteq N$.

Of course, $\mathrm{Mod}(T)$ is a thin potentialist system. Note also that the structure of $\mathrm{Mod}(T)$ depends on the underlying domains of the models, in the sense that M may be a substructure of N but not of an isomorphic copy of N. It is natural then to consider allowing all embeddings between models, so that the relations between models depend only on their isomorphism types.

Definition 3.12. The *embedding potentialist system*, $\mathrm{ModEm}(T)$, based on T consists of the class of models of T together with all embeddings. In other words, from M to N we include all maps $M \to N$ which are embeddings of L-structures.

In Theorem 3.28 below, I will show that $\mathrm{Mod}(T)$ and $\mathrm{ModEm}(T)$ are bitotally bisimilar in a strong sense. As noted in the introduction, this provides and important and deep connection between the philosophically interesting $\mathrm{Mod}(T)$ on the one hand, and the mathematically tractable $\mathrm{ModEm}(T)$ on the other.

3.3 Bisimulations

The notion of bisimulation is central in propositional modal logic (see for example [BRV01, §2.2]). It can be extended to first-order modal logic, at the same time generalising the notion of a back-and-forth system for first-order structures. The following is adapted from Definition 11.4.2 in [Ben10, p. 123].

Definition 3.13. A *bisimulation* between potentialist systems $\mathcal{M}$ and $\mathcal{N}$ in the same language L is a non-empty relation $\sim$ matching pairs $(M, \overline{a})$ with pairs $(N, \overline{b})$, where $(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$ are pointed systems of the same parameter size, such that when $(M, \overline{a}) \sim (N, \overline{b})$ the following hold.

(B1) The map $\overline{a} \mapsto \overline{b}$ is a partial isomorphism between M and N.

(B2) For any $c \in M$ there is $d \in N$ such that $(M, \overline{a}c) \sim (N, \overline{b}d)$, and vice versa (where $\overline{a}c$ is the result of adding c at then end of the tuple $\overline{a}$, etc.).

(B3) For any $\pi \colon M \to M'$ in $\mathcal{M}$ there is $\rho \colon N \to N'$ in $\mathcal{N}$ such that:

$$(M', \pi(\overline{a})) \sim (N', \rho(\overline{b}))$$

and vice versa.

Write $(\mathcal{M}, M, \overline{a}) \leftrightarroweq (\mathcal{N}, N, \overline{b})$ when there is a bisimulation $\sim$ between $\mathcal{M}$ and $\mathcal{N}$ such that:

$$(M, \overline{a}) \sim (N, \overline{b})$$

Definition 3.14. A bisimulation $\sim$ is *left total* if every pair $(M, \overline{a})$ on the left is related to some pair $(N, \overline{b})$ on the right. It is *right total* if every pair $(N, \overline{b})$ on the right is related to some pair $(M, \overline{a})$ on the left. It is *bitotal* if it is both left and right total.

A bisimulation is a kind of back-and-forth system corresponding to a modal version of the Ehrenfeucht-Fraïssé game. This game-theoretic side will be examined in more detail in Section 3.4. There, it will also be shown that bisimulations correspond to $L_\infty^\Diamond$-elementary equivalence. We can get a more fine-grained notion corresponding to $L_\infty^{\Diamond,\alpha+1}$-elementary equivalence as follows.

Definition 3.15. Let α be an ordinal. An *α-bisimulation* between potentialist systems $\mathcal{M}$ and $\mathcal{N}$ in the same language L is a collection of non-empty relations $\sim_\beta$ for $\beta \leqslant \alpha$. Each relates pairs $(M, \overline{a})$ with pairs $(N, \overline{b})$, where $(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$ are pointed systems of the same parameter size. Whenever $(M, \overline{a}) \sim_\beta (N, \overline{b})$ we require that the following hold.

(αB1) The map $\overline{a} \mapsto \overline{b}$ is a partial isomorphism between M and N.

(αB2) For every $\gamma < \beta$: for any $c \in M$ there is $d \in N$ such that we have $(M, \overline{a}c) \sim_\gamma (N, \overline{b}d)$, and vice versa.

(αB3) For every $\gamma < \beta$: for any $\pi\colon M \to M'$ in $\mathcal{M}$ there is $\rho\colon N \to N'$ in $\mathcal{N}$ such that $(M', \pi(\overline{a})) \sim_\gamma (N', \rho(\overline{b}))$, and vice versa.

Write $(\mathcal{M}, M, \overline{a}) \leftrightarroweq_\alpha (\mathcal{N}, N, \overline{b})$ to express the existence of an α-bisimulation between $\mathcal{M}$ and $\mathcal{N}$ such that $(M, \overline{a}) \sim_\alpha (N, \overline{b})$.

Remark 3.16. The reader might be wondering why an α-bisimulation is defined as a collection of relations $\sim_\beta$ for $\beta \leqslant \alpha$ rather than for $\beta < \alpha$. It would appear that we lack a convenient way of expressing that '$(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$ are α-bisimilar for every α less than some limit ordinal δ'. However, it is not hard to see that this notion just described is equivalent to $(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$ being δ-bisimilar. Moreover, it is desirable that being 0-bisimilar should coincide with being equivalent on atomic formulas.

Remark 3.17. In Definition 3.15 there is no requirement made that the collection of relations $\sim_\beta$ 'cohere' in any sense. But note that without loss of generality we can always assume that $\sim_\beta \subseteq \sim_\gamma$ for $\gamma < \beta$.

Now, one way of thinking about a bisimulation is as a collection of partial isomorphisms of finite parts of first-order structures which can extended indefinitely. A special case of this is when these partial isomorphisms uniquely combine to a system of total isomorphisms between worlds in one potentialist system and worlds in the other.

Definition 3.18. An *iso-bisimulation* between potentialist systems $\mathcal{M}$ and $\mathcal{N}$ is a non-empty relation $\approx$ between $\mathcal{M}$ and $\mathcal{N}$ together with a system of isomorphisms $(\xi_{M,N}\colon M \to N \mid M \approx N)$ subject to the following condition. Whenever $\pi\colon M \to M'$ in $\mathcal{M}$ and $M \approx N$, there is $\rho\colon N \to N'$ in $\mathcal{N}$ such that $M' \approx N'$ and the following diagram commutes; and vice versa.

$$\begin{array}{ccc} M' & \xrightarrow{\xi_{M',N'}} & N' \\ {\scriptstyle\pi}\uparrow & & \uparrow{\scriptstyle\rho} \\ M & \xrightarrow{\xi_{M,N}} & N \end{array}$$

The following allows us to convert any iso-bisimulation into an (ordinary) bisimulation.

Lemma 3.19. *Let* $(\approx, (\xi))$ *be an iso-bisimulation between* $\mathcal{M}$ *and* $\mathcal{N}$. *Define the relation* $\approx^*$ *by setting* $(M, \overline{a}) \approx^* (N, \overline{b})$ *if and only if* $M \approx N$ *and* $\xi_{M,N}(\overline{a}) = \overline{b}$. *Then* $\approx^*$ *is a bisimulation.*

Hence, any iso-bisimulation gives rise to a bisimulation. The converse is not true in general, e.g. for cardinality reasons, or because a bisimulation can involve multiple non-compatible partial isomorphisms between structures. The following is a simple example of two bisimilar systems which are not iso-bisimilar.

Example 3.20. Let L be the empty language. Let $\mathcal{M}$ be the potentialist system consisting of a single $\aleph_0$-size L-structure, and let $\mathcal{N}$ be the structure consisting of an $\aleph_0$-sized and an $\aleph_1$-sized L-structure, with an embedding from the former to the latter. There is no iso-bisimulation $\mathcal{M} \leftrightarroweq \mathcal{N}$ for cardinality reasons, but there is an (ordinary) bisimulation $\mathcal{M} \leftrightarroweq \mathcal{N}$. Indeed, relate pairs according to their atomic types:[1]

$$(M, \overline{a}) \sim (N, \overline{b}) \quad \Leftrightarrow \quad \Delta_0^M(\overline{a}) = \Delta_0^N(\overline{b})$$

A key property of bisimulations is their invariance: bisimilar systems satisfy the same formulas (c.f. [Ben10, Theorem 21, p. 124]).

Theorem 3.21 (Fundamental Theorem of Bisimulations)**.** *If two pointed systems are bisimilar then they are* $L_\infty^\Diamond$*-elementarily equivalent.*

Proof. The proof is by induction on formulas. Let $\sim$ be a bisimulation between $(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$. The atomic case follows from (B1). If we have $\mathcal{M}, M \vDash \exists x \psi([\overline{a}], x)$ then there is $c \in M$ such that $\mathcal{M}, M \vDash \psi[\overline{a}, c]$; by (B2) there is $d \in N$ such that $(M, \overline{a}c) \sim (N, \overline{b}d)$, so by induction hypothesis $\mathcal{N}, N \vDash \psi[\overline{b}, d]$ and $\mathcal{N}, N \vDash \exists x \psi([\overline{b}], x)$. If $\mathcal{M}, M \vDash \Diamond \psi[\overline{a}]$ then for some $\pi \colon M \to M'$ in $\mathcal{M}$ we have $\mathcal{M}, M' \vDash \psi[\pi(\overline{a})]$. By (B3) there is $\rho \colon N \to N'$ in $\mathcal{N}$ with $(M', \pi(\overline{a})) \sim (N', \rho(\overline{b}))$; so by induction hypothesis $\mathcal{N}, N' \vDash \psi[\rho(\overline{b})]$ and $\mathcal{N}, N \vDash \Diamond \psi[\overline{b}]$. □

Corollary 3.22. *If* $(\approx, (\xi))$ *is an iso-bisimulation between* $\mathcal{M}$ *and* $\mathcal{N}$*, then whenever* $M \approx N$*, for any* $\overline{a}$ *in* M *and formula* ϕ *of* $L_\infty^\Diamond$*:*

$$\mathcal{M}, M \vDash \phi[\overline{a}] \Leftrightarrow \mathcal{N}, N \vDash \phi[\xi_{M,N}(\overline{a})]$$

Theorem 3.23 (Fundamental Theorem of α-bisimulations)**.** *If two pointed systems are* α*-bisimilar then they are* $L_\infty^{\Diamond;\alpha+1}$*-elementarily equivalent.*

Proof. This is a refinement of the proof of Theorem 3.21. We prove by induction on modal rank that whenever $(M, \overline{a}) \sim_\beta (N, \overline{b})$, for any ϕ with $\mathrm{mqr}(\phi) \leqslant \beta$ we have:

$$\mathcal{M}, M \vDash \phi[\overline{a}] \quad \Leftrightarrow \quad \mathcal{N}, N \vDash \phi[\overline{b}]$$

[1] I use the notation $\Delta_0^M(\overline{a})$ to denote the atomic type of $\overline{a}$ in M: the set of atomic formulas true in M of $\overline{a}$.

As above the base case is dealt with by (αB1). If $M \vDash \exists x \psi([\overline{a}], x)$ there is $c \in M$ such that $M \vDash \psi[\overline{a}, c]$; by ($\alpha$B2) there is $d \in N$ such that:

$$(M, \overline{a}c) \sim_{\mathrm{mqr}(\psi)} (N, \overline{b}d)$$

which by induction hypothesis means that $N \vDash \psi[\overline{b}, d]$. The other cases are similar. □

The converses of Theorem 3.21 and Theorem 3.23 will be proved in Section 3.4 after we have established the connection with infinite games.

The following provide promised proofs of facts mentioned in Section 3.2. We first see that any potentialist system is iso-bisimilar with a thin one, using the 'unravelling' construction, a generalisation from propositional modal logic (see [BRV01, Definition 4.51]).

Definition 3.24. Let $\mathcal{M}$ be a potentialist system. The *unravelling*, $\mathrm{Unrav}(\mathcal{M})$, of $\mathcal{M}$ is the free category on $\mathcal{M}$, regarded as a potentialist system. Concretely, $\mathrm{Unrav}(\mathcal{M})$ is specified as follows.

- The objects are finite sequences of embeddings:

$$\vec{\pi} = \left(M_0 \xrightarrow{\pi_1} M_1 \xrightarrow{\pi_2} \cdots \xrightarrow{\pi_{n-1}} M_{n-1} \xrightarrow{\pi_n} M_n \right)$$

  where $M_0, \ldots, M_n$ are worlds from $\mathcal{M}$ and $\pi_1, \ldots, \pi_n$ are embeddings from $\mathcal{M}$. We allow singleton sequences. The structure corresponding to $\vec{\pi}$ is $M_{\vec{\pi}} = M_n$.

- For every sequence $\vec{\pi}$ and every embedding π_{n+1} with domain $M_{\vec{\pi}}$, we obtain a new sequence $\vec{\pi}\pi_{n+1}$ by adding π_{n+1} at the end, and we have an arrow $\pi_{n+1} \colon \vec{\pi} \to \vec{\pi}\pi_{n+1}$.

Remark 3.25. This definition is slightly informal. According to Definition 3.1, worlds in a potentialist system should be L-structures, but here they are finite sequences. So to formally define the $\mathrm{Unrav}(\mathcal{M})$ we should take the worlds to be the last structures $M_{\vec{\pi}} = M_n$, and somehow encode the sequence structure in the domain. For example, we could rename the elements of the domain as $M_{\vec{\pi}} \times \{(\pi_1 \cdots \pi_n)\}$. I won't worry too much about these technicalities from now on.

Note that the unravelling of $\mathcal{M}$ is always a thin category. The following gives the bisimulation $\mathcal{M} \leftrightarroweq \mathrm{Unrav}(\mathcal{M})$.

Lemma 3.26. *Let $\mathcal{M}$ be a potentialist system. The relation $\approx$ between $\mathcal{M}$ and* $\mathrm{Unrav}(\mathcal{M})$, *given by: $N \approx \vec{\pi}$ if and only if $M_{\vec{\pi}} = N$, together with the identity isomorphisms, is an iso-bisimulation which is bitotal.*

Proof. Assume that $N \approx \vec{\pi}$. Take $\rho\colon N \to N'$ in $\mathcal{M}$. Then $\vec{\pi}\rho \in \mathrm{Unrav}(\mathcal{M})$ and $N' \approx \vec{\pi}\rho$. Furthermore, we have:

$$\rho\colon \vec{\pi} \to \vec{\pi}\rho$$

in $\mathrm{Unrav}(\mathcal{M})$. Conversely, if $\rho\colon \vec{\pi} \to \vec{\pi}\rho$ in $\mathrm{Unrav}(\mathcal{M})$ with codomain N', then we have $\rho\colon N \to N'$ in $\mathcal{M}$. □

Corollary 3.27. *Every potentialist system is bisimilar to thin one via a bitotal iso-bisimulation.*

Let us see now that $\mathrm{Mod}(T)$ and $\mathrm{ModEm}(T)$ are bisimilar.

Theorem 3.28. *Let T be a first-order L-theory. There is a bitotal iso-bisimulation between the systems* $\mathrm{Mod}(T)$ *and* $\mathrm{ModEm}(T)$.

We first need to define a basic disjointifying bisimulation, which makes sure that worlds are disjoint.

Definition 3.29. Let $\mathcal{M}$ be a potentialist system. The *disjointification* of $\mathcal{M}$, notation $\mathcal{M}^\bullet$, is the potentialist system obtained from $\mathcal{M}$ by replacing the domain of each world M with $M \times \{M\}$, then taking all the embeddings from $\mathcal{M}$, suitably modified. The *disjointifying bisimulation* $\curlywedge$ is the iso-bisimulation $\mathcal{M} \leftrightarroweq \mathcal{M}^\bullet$ defined by relating each world in $\mathcal{M}^\bullet$ with its original copy in $\mathcal{M}$:

$$M \curlywedge M \times \{M\}$$

Now, the idea to prove Theorem 3.28 is that we first take the unravelling $\mathrm{Unrav}(\mathrm{ModEm}(T)^\bullet)$, and then start renaming the elements of its worlds so that embeddings become inclusions. In order to do this consistently, we need to carry out this process iteratively through sequences of embeddings. See Figure 3.1 for an illustration. Fortunately, every world in $\mathrm{Unrav}(\mathrm{ModEm}(T)^\bullet)$ remembers a sequence of previous embeddings, and this allows us to carry out the process.

Definition 3.30. Let $\mathcal{M}$ be a potentialist system, and take a sequence $\vec{\pi} = \left(M_0 \xrightarrow{\pi_1} \cdots \xrightarrow{\pi_n} M_n\right) \in \mathrm{Unrav}(\mathcal{M}^\bullet)$. The structure $\mathrm{Coalesce}(\vec{\pi})$ is the isomorphic copy of M_n defined iteratively by renaming $\pi_1[M_0]$ as M_0 in M_1, then renaming $\pi_2[M_1]$ using the result of this, and so on, as illustrated in Figure 3.1.

Remark 3.31. The disjointified version $\mathcal{M}^\bullet$ of $\mathcal{M}$ is necessary because the structure $\mathrm{Coalesce}(\vec{\pi})$ is only sensible if no renamed element becomes the same as a non-renamed element. For instance, in the case $n = 1$, if there is $a \in M_0 \cap (M_1 \setminus \pi_1[M_0])$, then the elements $a, \pi_1(a) \in M_0$ will have collapsed to a single one $a \in \mathrm{Coalesce}(\vec{\pi})$.

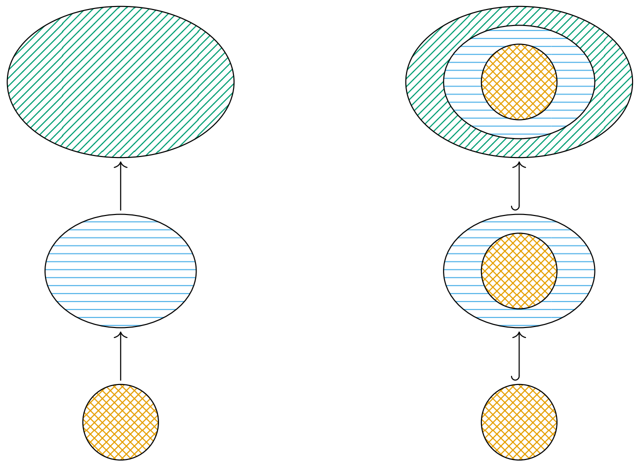

Figure 3.1: An example showing that iteratively renaming elements is necessary. Blobs represent worlds and arrows embeddings. Hook arrows represent inclusions. To transform the sequence on the left so that embeddings become inclusions, we need to iteratively rename along the embeddings, as on the right.

Definition 3.32. Let T be any first-order L-theory. Let the *renaming bisimulation* $\propto$ between $\mathrm{Mod}(T)^\bullet$ and $\mathrm{Unrav}(\mathrm{ModEm}(T)^\bullet)$ be the iso-bisimulation defined as follows. For any $\vec{\pi} \in \mathrm{Unrav}(\mathrm{ModEm}(T)^\bullet)$ put:

$$\mathrm{Coalesce}(\vec{\pi}) \propto \vec{\pi}$$

The isomorphism $\mathrm{Rename}_{\vec{\pi}} \colon \mathrm{Coalesce}(\vec{\pi}) \to M_{\vec{\pi}}$ is the natural renaming isomorphism.

Theorem 3.33. *The renaming bisimulation is a bitotal iso-bisimulation.*

Proof. First note that for any $M_0 \in \mathrm{Mod}(T)^\bullet$, we have a length-0 sequence $\epsilon_{M_0} \in \mathrm{Unrav}(\mathrm{ModEm}(T)^\bullet)$ with M_0 as the initial world, and $M_0 \propto \epsilon_{M_0}$; hence $\propto$ is left total. It is clearly right total. So we need to check the isomorphism extension property for iso-bisimulations. Take $\vec{\pi} \in \mathrm{Unrav}(\mathrm{Mod}(T)^\bullet)$, so that $\mathrm{Coalesce}(\vec{\pi}) \propto \vec{\pi}$.

Assume that $\mathrm{Coalesce}(\vec{\pi}) \subseteq N$ in $\mathrm{Mod}(T)$. Define $\rho \colon M_{\vec{\pi}} \to N$ to be the result of renaming in $M_{\vec{\pi}}$ via $\mathrm{Rename}_{\vec{\pi}}^{-1}$, then applying the inclusion embedding into N. Then in fact:

$$N = \mathrm{Coalesce}(\vec{\pi}\rho)$$

We have $\rho \colon \vec{\pi} \to \vec{\pi}\rho$ in $\mathrm{Unrav}(\mathrm{ModEm}(T)^\bullet)$, and the following commutative diagram, where the hook arrow represents an inclusion. See also Figure 3.2 for a picture-based commutative diagram.

$$
\begin{array}{ccc}
N = \mathrm{Coalesce}(\vec{\pi}\rho) & \xrightarrow{\mathrm{Rename}_{\vec{\pi}\rho}} & N \\
\big\uparrow & & \big\uparrow \rho \\
\mathrm{Coalesce}(\vec{\pi}) & \xrightarrow{\mathrm{Rename}_{\vec{\pi}}} & M_{\vec{\pi}}
\end{array}
$$

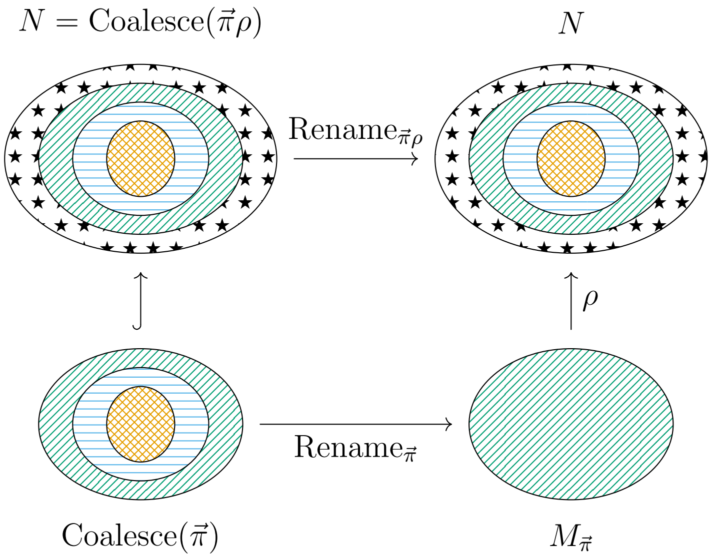

Figure 3.2: Defining $\rho\colon M_{\vec{\pi}} \to N$ from $\mathrm{Coalesce}(\vec{\pi}) \subseteq N$. The hook arrow represents an inclusion.

Conversely, take $\rho\colon \vec{\pi} \to \vec{\pi}\rho$ in $\mathrm{Unrav}(\mathrm{ModEm}(T)^{\bullet})$ with codomain N. Note that by definition $\mathrm{Coalesce}(\vec{\pi}) \subseteq \mathrm{Coalesce}(\vec{\pi}\rho)$, and this is a substructure relation. Moreover, the following diagram commutes.

$$
\begin{array}{ccc}
\mathrm{Coalesce}(\vec{\pi}\rho) & \xrightarrow{\mathrm{Rename}_{\vec{\pi}\rho}} & N \\
\big\uparrow & & \big\uparrow \rho \\
\mathrm{Coalesce}(\vec{\pi}) & \xrightarrow{\mathrm{Rename}_{\vec{\pi}}} & M_{\vec{\pi}}
\end{array}
$$

□

Proof of Theorem 3.28. Indeed:

$$\mathrm{Mod}(T) \leftrightarrows \mathrm{Mod}(T)^{\bullet} \leftrightarrows \mathrm{Unrav}(\mathrm{ModEm}(T)^{\bullet}) \leftrightarrows \mathrm{ModEm}(T)^{\bullet} \leftrightarrows \mathrm{ModEm}(T)$$

via bitotal iso-bisimulations. □

3.4 Modal Ehrenfeucht-Fraïssé games

We now turn to the connection between bisimulations and games. Modal Ehrenfeucht-Fraïssé games generalise their non-modal cousins (see [Hod93, §3.2]).

Definition 3.34. Let $(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$ be pointed systems of the same parameter size. The *modal Ehrenfeucht-Fraïssé game*:

$$\mathrm{MEF}(\mathcal{M}, M, \overline{a}, \mathcal{N}, N, \overline{b})$$

is played by Eloise and Abelard. The positions of the game are quadruples $(U, \overline{c}, Q, \overline{d})$, where $(\mathcal{M}, U, \overline{c})$ and $(\mathcal{N}, V, \overline{d})$ are pointed systems of the same parameter size. The initial position is $(M, \overline{a}, N, \overline{b})$. The game proceeds in moves alternating between Abelard and Eloise. From position $(U, \overline{c}, V, \overline{d})$, Abelard can make one of the following two kinds of moves.

(a) Choose $u \in U$ or $v \in V$.

(b) Choose $\pi\colon U \to U'$ in $\mathcal{M}$ or $\rho\colon V \to V'$ in $\mathcal{N}$.

Eloise then responds, respectively, as follows.

(a) When Abelard chose $u \in U$, choose $v \in V$; otherwise choose $u \in U$. The new position is then $(U, \overline{c}u, V, \overline{d}v)$.

(b) When Abelard chose $\pi\colon U \to U'$, choose $\rho\colon V \to V'$; otherwise choose $\pi\colon U \to U'$. The new position is then $(U', \pi(\overline{c}), V', \rho(\overline{d}))$.

A play in the game, consisting of an ω-sequence of positions, is a win for Eloise if and only if for every position $(U, \overline{c}, V, \overline{d})$ in the play, the atomic types agree (in other words $\overline{c} \mapsto \overline{d}$ is a partial isomorphism):

$$\Delta_0^U(\overline{c}) = \Delta_0^V(\overline{d})$$

Often, we want to consider just the set of all positions in the game based on $\mathcal{M}$ and $\mathcal{M}'$, together with the winning conditions, without fixing an initial position. I will call this the *board* and denote it by $\mathrm{MEF}(\mathcal{M}, \mathcal{M}')$. This becomes a game when we specify an initial position.

The following theorem demonstrates the connection between bisimulations and modal Ehrenfeucht-Fraïssé games. It simultaneously generalises the connection between non-modal Ehrenfeucht-Fraïssé games (based on two fixed structures) and back-and-forth systems, on the one hand [Hod93, Lemma 3.2.2], and bisimulation games and (propositional) bisimulations, on the other [Ben14, §15.7].

Theorem 3.35. *Eloise has a winning strategy in the game* $\mathrm{MEF}(\mathcal{M}, M, \overline{a}, \mathcal{N}, N, \overline{b})$ *if and only if* $(\mathcal{M}, M, \overline{a})$ *is bisimilar with* $(\mathcal{N}, N, \overline{b})$.

Proof. Let $\sim$ be a bisimulation $(\mathcal{M}, M, \overline{a}) \leftrightarroweq (\mathcal{N}, N, \overline{b})$. Eloise uses $\sim$ to see how she should play. Note that the conditions (B2) and (B3) guarantee that no matter what Abelard plays, Eloise can always ensure the play stays at positions $(U, \overline{c}, V, \overline{d})$ such that $(\mathcal{M}, U, \overline{c}) \sim (\mathcal{N}, V, \overline{d})$.[2] Then (B1) guarantees that for such positions:

$$\Delta_0^U(\overline{c}) = \Delta_0^V(\overline{d})$$

meaning that the resulting play is a win for Eloise.

Now assume that Eloise has a winning strategy in $\mathrm{MEF}(\mathcal{M}, M, \overline{a}, \mathcal{N}, N, \overline{b})$. We can construct a bisimulation by playing as Abelard and seeing what Eloise does. Put $(U, \overline{c}) \sim (V, \overline{d})$ if and only if the position $(U, \overline{c}, V, \overline{d})$ occurs in some play in which Eloise plays her winning strategy. Let us see that this is a bisimulation. Condition (B1) follows from the fact that the atomic types of such positions must agree. As for condition (B2), take $u \in U$. We make Abelard choose this $u \in U$, and then according to her winning strategy Eloise then plays some $v \in V$, with the result that $(U, \overline{c}u) \sim (V, \overline{d}v)$. The converse direction, and condition (B3), are similar. □

In the first-order case, the existence of a winning strategy for Eloise in the Ehrenfeucht-Fraïssé game can be expressed by a formula in an infinitary language (see [Hod93, §3.5]). An analogous result holds for modal Ehrenfeucht-Fraïssé games, allowing us to show that two pointed systems are bisimilar if and only if they are $L_\infty^\Diamond$-elementarily equivalent. This further yields a converse to the Fundamental Theorem of Bisimulations (Theorem 3.21).

Theorem 3.36. *The following are equivalent.*

(1) Eloise has a winning strategy in $\mathrm{MEF}(\mathcal{M}, M, \overline{a}, \mathcal{N}, N, \overline{b})$.

(2) $(\mathcal{M}, M, \overline{a})$ *is* $L_\infty^\Diamond$*-elementarily equivalent to* $(\mathcal{N}, N, \overline{b})$.

(3) There is a bisimulation $(\mathcal{M}, M, \overline{a}) \leftrightarroweq (\mathcal{N}, N, \overline{b})$.

The version of the equivalence (1) ⇔ (2) for non-modal Ehrenfeucht-Fraïssé games is Theorem 3.5.2 in [Hod93]. I follow the proof structure given there, modifying to account for: (1) the additional modal aspect, and (2) the fact that potentialist systems can be class-sized. We first need some terminology concerning the ranks of game positions.

[2] Note that in general we need global choice in order to define Eloise's strategy from the bisimulation.

Definition 3.37. The *rank*, notation $\operatorname{rank}(p)$, of a position $p = (M, \overline{a}, N, \overline{b})$ in $\mathrm{MEF}(\mathcal{M}, \mathcal{N})$ is either -1, an ordinal, or ∞, and is determined as follows.

- $\operatorname{rank}(p) \geqslant 0$ if and only if $\Delta_0^M(\overline{a}) = \Delta_0^N(\overline{a})$.
- $\operatorname{rank}(p) \geqslant \alpha + 1$ if and only if for every possible move made by Abelard from p, Eloise can then move to a position of rank at least α.
- $\operatorname{rank}(p) \geqslant \lambda$, for λ a limit ordinal, if and only if $\operatorname{rank}(p) \geqslant \alpha$ for all $\alpha < \lambda$.

Lemma 3.38. *For any ordinal α and pointed systems $(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$ of the same parameter size we have:*

$$\operatorname{rank}(M, \overline{a}, N, \overline{b}) \geqslant \alpha \text{ in } \mathrm{MEF}(\mathcal{M}, \mathcal{N}) \quad \Leftrightarrow \quad (\mathcal{M}, M, \overline{a}) \leftrightarroweq_\alpha (\mathcal{N}, N, \overline{b})$$

Proof. This is a refinement of the proof of Theorem 3.35. Put $(U, \overline{c}) \sim_\beta (V, \overline{d})$ if and only if $\beta \leqslant \min\{\operatorname{rank}(U, \overline{c}, V, \overline{d}), \alpha\}$. This then gives an α-bisimulation $(\mathcal{M}, M, \overline{a}) \leftrightarroweq_\alpha (\mathcal{N}, N, \overline{b})$. Conversely, if if we have any α-bisimulation $(\mathcal{M}, M, \overline{a}) \leftrightarroweq_\alpha (\mathcal{N}, N, \overline{b})$, then any instance $(U, \overline{c}) \sim_\beta (V, \overline{d})$ witnesses that $\beta \leqslant \operatorname{rank}(U, \overline{c}, V, \overline{d})$. □

Lemma 3.39.

(1) If α is an ordinal such that every position of rank at least α is also at least $\alpha + 1$, then every position of rank at least α is winning for Eloise.

(2) A position has rank ∞ if and only if it is winning for Eloise. Hence:

$$(\mathcal{M}, M, \overline{a}) \leftrightarroweq (\mathcal{N}, N, \overline{b}) \quad \Leftrightarrow \quad \forall \alpha \colon (\mathcal{M}, M, \overline{a}) \leftrightarroweq_\alpha (\mathcal{N}, N, \overline{b})$$

Proof. These facts are proved more generally in Lemma 3.4.1 and Theorem 3.4.2 of [Hod93]. □

As mentioned above, we can express the property of having rank at least α using an $L_\infty^\Diamond$-formula of modal-quantifier rank α. I will omit the $\mathcal{M}$ from the superscript of θ whenever possible.

Theorem 3.40. *Let $(\mathcal{M}, M, \overline{a})$ be a pointed system of parameter size n. Take an ordinal α. There is an n-variable formula $\theta_\alpha^{\mathcal{M}, M, \overline{a}}$ in $L_\infty^\Diamond$ of modal-quantifier rank exactly α such that for every pointed system $(\mathcal{N}, N, \overline{b})$ of parameter size n, we have:*

$$\mathcal{N}, N \vDash \theta_\alpha^{\mathcal{M}, M, \overline{a}} \left[\overline{b}\right] \quad \Leftrightarrow \quad \operatorname{rank}(M, \overline{a}, N, \overline{b}) \geqslant \alpha \text{ in } \mathrm{MEF}(\mathcal{M}, \mathcal{N})$$

Proof. The formula $\theta_\alpha^{M, \overline{a}}$ is defined by recursion on α.

- First:

$$\theta_0^{M,\overline{a}}(\overline{x}) := \bigwedge \Delta_0^M(\overline{a})$$

- For successor ordinals $\alpha + 1$, we need to make sure that Eloise can always force to a position of rank at least α, no matter what Abelard chooses. Considering the definition of the game, this requires expressing the following.

  (i) For every $u \in M$ there is $y \in N$ such that $\theta_\alpha^{M,(\overline{a}u)}(\overline{x}, y)$ holds at N.

  (ii) For every $y \in N$ there is $u \in M$ such that $\theta_\alpha^{M,(\overline{a}u)}(\overline{x}, y)$ holds at N.

  (iii) For every $\pi \colon M \to M'$ there is $\rho \colon N \to N'$ such that $\theta_\alpha^{M,\pi(\overline{a})}(\rho(\overline{x}))$ holds at N'.

  (iv) For every $\rho \colon N \to N'$ there is $\pi \colon M \to M'$ such that $\theta_\alpha^{M,\pi(\overline{a})}(\rho(\overline{x}))$ holds at N'.

  The first two conditions can be expressed by replacing quantification over M with conjunctions and disjunctions. For the other two conditions, this strategy does not immediately work, since there may be class-many embeddings $\pi \colon M \to M'$, but we are only allowed to take set-sized conjunctions and disjunctions. Fortunately, it will follow from the induction that there are only set-many formulas of the form $\theta_\alpha^{M,\pi(\overline{a})}(\overline{x})$.[3] Let Φ be the set of formulas $\theta_\alpha^{M,\pi(\overline{a})}(\overline{x})$ for $\pi \colon M \to M'$ (Even though there may be class-many *names* $\theta_\alpha^{M,\pi(\overline{a})}$, there are only set-many formulas named by them.). We can now define:

$$\theta_{\alpha+1}^{M,\overline{a}}(\overline{x}) := \left(\begin{array}{ll} & \bigwedge_{u \in M} \exists y \theta_\alpha^{M,(\overline{a}u)}(\overline{x}, y) \\ \wedge & \forall y \bigvee_{u \in M} \theta_\alpha^{M,(\overline{a}u)}(\overline{x}, y) \\ \wedge & \bigwedge_{\phi \in \Phi} \Diamond \phi \\ \wedge & \Box \bigvee_{\phi \in \Phi} \phi \end{array} \right)$$

- Finally, when γ is a limit, we define:

$$\theta_\gamma^{M,\overline{a}}(\overline{x}) := \bigwedge_{\alpha < \gamma} \theta_\alpha^{M,\overline{a}}(\overline{x})$$

Note that at each stage the definition of $\theta_\alpha^{M,\overline{a}}$ corresponds with the condition that $\mathrm{rank}(M, \overline{a}, N, \overline{b}) \geqslant \alpha$. □

With this key piece we can establish a correspondence between all of the notions we have of 'equivalence up to rank α'.

[3] In fact, one can show that there are only set-many formulas of rank α up to equivalence.

Theorem 3.41. *Let $(\mathcal{M}, M, \overline{a})$ and $(\mathcal{N}, N, \overline{b})$ be pointed systems of the same parameter size, and let α be an ordinal. The following are equivalent.*

(1) $\mathcal{M}, M \vDash \theta_\alpha^{\mathcal{N},N,\overline{b}}[\overline{a}]$.

(2) $\operatorname{rank}(M, \overline{a}, N, \overline{b}) \geqslant \alpha$ *in* $\operatorname{MEF}(\mathcal{M}, \mathcal{N})$.

(3) $(\mathcal{M}, M, \overline{a}) \Leftrightarrow_\alpha (\mathcal{N}, N, \overline{b})$.

(4) $(\mathcal{M}, M, \overline{a}) \equiv_\infty^{\Diamond,\alpha+1} (\mathcal{N}, N, \overline{b})$.

Proof. (1) $\Leftrightarrow$ (2) is by Theorem 3.40, (2) $\Leftrightarrow$ (3) is by Lemma 3.38, and (3) $\Rightarrow$ (4) is by Theorem 3.23.

(4) $\Rightarrow$ (1). Note that $(N, \overline{b}, N, \overline{b})$ is winning for Eloise in $\operatorname{MEF}(\mathcal{N}, \mathcal{N})$, and hence has rank at least α by (2) of Lemma 3.39. Hence by Theorem 3.40 we have:

$$\mathcal{N}, N \vDash \theta_\alpha^{\mathcal{N},N,\overline{b}}\left[\overline{b}\right]$$

But $\theta_\alpha^{\mathcal{N},N,\overline{b}}$ has modal-quantifier rank α, whence by assumption:

$$\mathcal{M}, M \vDash \theta_\alpha^{\mathcal{N},N,\overline{b}}[\overline{a}]$$ □

Putting everything together, we can now prove the main result.

Proof of Theorem 3.36. We have:

$$\begin{aligned}
M, \overline{a} \equiv_\infty^{\Diamond} N, \overline{b} \quad &\Leftrightarrow \quad \forall \alpha\colon M, \overline{a} \equiv_\infty^{\Diamond,\alpha+1} N, \overline{b} \\
&\Leftrightarrow \quad \forall \alpha\colon (\mathcal{M}, M, \overline{a}) \Leftrightarrow_\alpha (\mathcal{N}, N, \overline{b}) && \text{(Theorem 3.41)} \\
&\Leftrightarrow \quad \text{Eloise will win } \operatorname{MEF}(\mathcal{M}, M, \overline{a}, \mathcal{N}, N, \overline{b}) && \text{(Lemma 3.39)} \\
&\Leftrightarrow \quad (\mathcal{M}, M, \overline{a}) \Leftrightarrow (\mathcal{N}, N, \overline{b}) && \text{(Theorem 3.35)}
\end{aligned}$$

□

3.5 Mod(T) and set-sized potentialist systems

The remainder of this chapter will consider following the question.

Question 3.42. *When is a potentialist system bitotally bisimilar to one consisting of set-many worlds?*

In this section, I focus on those systems of the form Mod(T) for some first-order L-theory T, and ask the following specialisation of this question.

Question 3.43. *For which first-order theories T is* $\mathrm{Mod}(T)$ *bitotally bisimilar a set-sized potentialist system?*

In this section and the next I make use of basic model theoretic terminology and techniques, such as $\aleph_0$-saturated models and the method of diagrams. For a reference see [Hod93].

Let us first see that answer to Question 3.43 is not "all of them", by considering the theory of directed graphs. Indeed, this potentialist system has a very rich structure; see [HW20] for further details.

Proposition 3.44. *Let $L = \{R\}$ and let T be the L-theory of directed graphs. Then* $\mathrm{Mod}(T)$ *is not bitotally bisimilar with a set-sized system.*

Proof. Using the Fundamental Theorem of Bisimulations (Theorem 3.21), it suffices to show that there is a proper class of directed graphs which are pair-wise distinguishable using $L^{\Diamond}_{\infty}$-formulas. In fact we don't need to use any modalities for this.

For every ordinal α the set $(\alpha+1, \in)$ is a model of T. The structure of each ordinal can be encoded in an L_∞-sentence as follows. Define the formula $\xi_\alpha(x)$ inductively by:

$$\xi_\alpha(x) := \bigwedge_{\beta<\alpha} (\exists y(yRx \wedge \xi_\beta(y))) \wedge \forall y \left(yRx \rightarrow \bigvee_{\beta<\alpha} \xi_\beta(y) \right)$$

It can be shown by induction that $(\alpha + 1, \in) \vDash \xi_\beta[\alpha]$ if and only if $\alpha = \beta$. Now let:

$$\nu_\alpha := \exists x(\forall y \neg xRy \wedge \xi_\alpha(x))$$

Since every structure $(\alpha + 1, \in)$ has a unique top element α, we then have that $(\alpha + 1, \in) \vDash \nu_\beta$ if and only if $\alpha = \beta$. □

To provide a positive answer to Question 3.43, we seek a property of first-order theories which ensures that $\mathrm{Mod}(T)$ is not 'unboundedly complicated'. The following new notion provides a sufficient condition.

Definition 3.45. Say that a theory T is *κ-rich* for some cardinal $\kappa \geqslant |L| + \aleph_0$ if it has infinite models, and all its models of size κ are $\aleph_0$-saturated.

Let us first see a couple of basic results about this notion.

Lemma 3.46. *If T is κ-rich, and $\lambda \geqslant \kappa$, then T is also λ-rich.*

Proof. Indeed, let $M \models T$ be a λ-model. Take $S \subseteq M$ finite. By the Downward Löwenheim-Skolem Theorem, there is an elementary substructure $N \preceq M$ of size κ, containing S, which is a model of T. Since T is κ-rich, every type over S is realised in N, and hence too in M. □

Lemma 3.47. *Let T be such that every completion is κ-categorical. Then T is κ-rich.*

Proof. If all completions are κ-categorical, then all κ-models are saturated, so in particular $\aleph_0$-saturated. □

Note that the converse doesn't hold for $\kappa > |L|$. For example, if T is $|L|$-categorical then it is $|L|^+$-rich, but it need not be $|L|^+$-categorical.

We will see that when T is κ-rich $\mathrm{Mod}(T)$ is bisimilar with a certain set-sized system. This system is defined by removing all the worlds from $\mathrm{Mod}(T)$ of size greater than κ, as follows.

Definition 3.48. Let λ be a cardinal. Define $\mathrm{Mod}_\lambda(T)$ to be the subsystem of $\mathrm{Mod}(T)$ consisting of those models of cardinality less than λ. The system $\mathrm{ModEm}_\lambda(T)$ is defined similarly.

By restricting the relation, the renaming bisimulation readily yields a bisimulation between the cardinality-restricted systems.

Lemma 3.49. *The renaming bisimulation* $\propto\colon \mathrm{Mod}(T) \leftrightarroweq \mathrm{ModEm}(T)$ *restricts to a bitotal iso-bisimulation:*

$$\mathrm{Mod}_\lambda(T) \leftrightarroweq \mathrm{ModEm}_\lambda(T)$$

Finally, potentialist systems such as $\mathrm{Mod}(T)$ and $\mathrm{Mod}_\lambda(T)$ contain class-many isomorphic copies of each model. Since we are interested in reducing the size of such systems, it will be necessary to consider the equivalent system containing only one model for each isomorphism type.

Definition 3.50. Let $\mathcal{M}$ be a potentialist system. Its *skeleton*, $\mathrm{Skel}(\mathcal{M})$, is the skeleton in the category-theoretic sense, i.e. the quotient under the relation which identifies M and M' if there are mutually inverse embeddings between them in $\mathcal{M}$. Concretely, for each equivalence class, we pick one representative, and let $\mathrm{Skel}(\mathcal{M})$ be the subsystem of $\mathcal{M}$ consisting of these worlds and all embeddings between them.

Remark 3.51. Giving $\mathrm{Skel}(\mathcal{M})$ a concrete realisation in general requires some class-level choice, like the Axiom of Global Choice. This can be avoided here, since we will only consider $\mathrm{Skel}(\mathrm{Mod}_\lambda(T))$ for some T, which can be realised by considering a set models of T whose domain is a subset of some fixed set of size λ.

Lemma 3.52. *The relation between $\mathcal{M}$ and* $\mathrm{Skel}(\mathcal{M})$ *which associates each structure to its equivalence class is a bitotal iso-bisimulation.*

Proof. This is immediate from the definition of iso-bisimulation. □

With the preliminaries in place, the main theorem for this section is ready to be proved.

Theorem 3.53. *If T is κ-rich then* $\mathrm{Mod}(T)$ *is bitotally bisimilar with* $\mathrm{Mod}_{\kappa^+}(T)$*, and thus with a set-sized potentialist system.*

Remark 3.54. Note that $\mathrm{Mod}(T)$ is not *iso*-bisimilar with a set-sized system, for cardinality reasons.

Proof of Theorem 3.53. Define $\sim$ between $\mathrm{ModEm}(T)$ and $\mathrm{ModEm}_{\kappa^+}(T)$ as follows. Take $(M, \overline{a})$ from $\mathrm{ModEm}(T)$ and $(N, \overline{b})$ from $\mathrm{ModEm}_{\kappa^+}(T)$. There are two cases.

- **Case 1:** $|M| \leqslant \kappa$. Put $(M, \overline{a}) \sim (N, \overline{b})$ if and only if $M, \overline{a} \cong N, \overline{b}$.
- **Case 2:** $|M| > \kappa$. Put $(M, \overline{a}) \sim (N, \overline{b})$ if and only if $|N| = \kappa$ and $\overline{a}$ and $\overline{b}$ have the same L-types: $\mathrm{Tp}_M(\overline{a}) = \mathrm{Tp}_N(\overline{b})$. For each $(M, \overline{a})$, such a $(N, \overline{b})$ must exist in $\mathrm{ModEm}_{\kappa^+}(T)$ by the Downward Löwenheim-Skolem Theorem.

Note that $\sim$ is bitotal. Let us verify that it is a bisimulation. Assume that $(M, \overline{a}) \sim (N, \overline{b})$. First note that, by definition, $M, \overline{a} \equiv N, \overline{b}$. Let us now consider condition (B2) on bisimulations. There are two cases.

Condition (B2), case 1: $|M| \leqslant \kappa$. Then $M, \overline{a} \cong N, \overline{b}$. Hence for each $c \in M$, there is $d \in N$ such that $(M, \overline{a}c) \sim (N, \overline{b}d)$, and conversely.

Condition (B2), case 2: $|M| > \kappa$. Then $\mathrm{Tp}_M(\overline{a}) = \mathrm{Tp}_N(\overline{b})$. By assumption on T and Lemma 3.46, both M and N are $\aleph_0$-saturated. Hence for any $c \in M$ there is $d \in N$ such that $\mathrm{Tp}_N(\overline{b}d) = \mathrm{Tp}_M(\overline{a}c)$, and conversely.

For condition (B3), there are again two cases.

Condition (B3), case 1: $|M| \leqslant \kappa$. Then by definition there is an isomorphism $f\colon (M, \overline{a}) \to (N, \overline{b})$. Assume that there is $\pi\colon M \to M'$ in $\mathrm{ModEm}(T)$. If $|M'| \leqslant \kappa$, then by definition there is $\rho\colon N \to N'$ in $\mathrm{ModEm}_{\kappa^+}(T)$ such that f extends to an isomorphism $M' \cong N'$. If, on the other hand, $|M'| > \kappa$, then by the Downward Löwenheim-Skolem Theorem, there is $N' \preceq M'$ of cardinality κ with $\pi(M) \subseteq N'$; note that $N' \in \mathrm{ModEm}_{\kappa^+}(T)$. Then we get an embedding $\rho := \pi \circ f^{-1}\colon N \to N'$, with the property that:

$$\mathrm{Tp}_{N'}(\rho(\overline{b})) = \mathrm{Tp}_{N'}(\pi(\overline{a})) = \mathrm{Tp}_{M'}(\pi(\overline{a}))$$

Whence $(M', \pi(\overline{a})) \sim (N', \rho(\overline{b}))$. See Figure 3.3 for a commutative diagram representing the situation. Conversely, any morphism $\rho\colon N \to N'$ in $\mathrm{ModEm}_{\kappa^+}(T)$ immediately gives a morphism $\pi\colon M \to M'$, which is the same up to isomorphism.

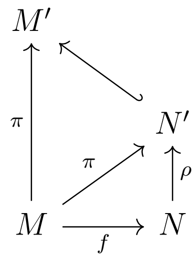

Figure 3.3: The commutative diagram for (B3), case 1

Condition (B3), case 2: $|M| > \kappa$. Take $\pi\colon M \to M'$ in $\mathrm{ModEm}(T)$. By a compactness argument, we can see that that the following is satisfiable (substituting the constants for $\overline{b}$ from $\mathrm{Diag}(N)$ for the free variables in $\mathrm{Tp}_{M'}(\pi(\overline{a}))$).

$$T \cup \mathrm{Diag}(N) \cup \mathrm{Tp}_{M'}(\pi(\overline{a}))(\overline{b})$$

Therefore, by the method of diagrams, there is a structure $N' \vDash T$ of cardinality κ with an embedding $\rho\colon N \to N'$ such that:

$$\mathrm{Tp}_{N'}(\rho(\overline{b})) = \mathrm{Tp}_{M'}(\pi(\overline{a}))$$

Whence $(M', \pi(\overline{a})) \sim (N', \rho(\overline{b}))$. Conversely, given any morphism $\rho\colon N \to N'$ in $\mathrm{ModEm}_{\kappa^+}(T)$ we can find a corresponding morphism $\pi\colon M \to M'$ by a very similar argument.

Finally, by Theorem 3.33 and Lemma 3.52, we have bitotal bisimulations:

$$\begin{aligned}\mathrm{Mod}(T) &\leftrightarroweq \mathrm{ModEm}(T)\\ &\leftrightarroweq \mathrm{ModEm}_{\kappa^+}(T)\\ &\leftrightarroweq \mathrm{Mod}_{\kappa^+}(T)\\ &\leftrightarroweq \mathrm{Skel}(\mathrm{Mod}_{\kappa^+}(T))\end{aligned}$$

the latter of which consists of set-many models. □

Corollary 3.55. *Let T be κ-categorical. Then* $\mathrm{Mod}(T)$ *is bitotally bisimilar with a system consisting of models of size $<\kappa$ and a single κ-model.*

Proof. Note that $\mathrm{Skel}(\mathrm{Mod}_{\kappa^+}(T))$ contains a single κ-size model. □

3.6 $\aleph_0$-categorical theories

In this section, I investigate Question 3.42 from a different direction. Given a potentialist system $\mathcal{M}$, we can consider its *theory* $\mathrm{Th}(\mathcal{M})$: the set of (finitary and non-modal) L-sentences which hold in all of its worlds. I first show that when $\mathrm{Th}(\mathcal{M})$ is $\aleph_0$-categorical and has quantifier elimination, then $\mathcal{M}$, irrespective its modal structure, is bisimilar to a very small potentialist system. On the other hand, I give an example showing that the result no longer holds when we drop the condition 'has quantifier elimination'.

Theorem 3.56. *Let $\mathcal{M}$ be a potentialist system which contains an infinite world, and let $T = \mathrm{Th}(\mathcal{M})$. Assume that T is $\aleph_0$-categorical and has quantifier elimination. Then $\mathcal{M}$ is bitotally bisimilar to a system which contains exactly one infinite world, which is countable, and which has no self-embeddings except the identity.*

In particular, when T has no finite models, $\mathcal{M}$ is bitotally bisimilar to a singleton system consisting of a countable world with just the identity embedding.

Remark 3.57. In Theorem 16 of [HW20] it is shown that if a theory T admits quantifier elimination, then it admits 'modality elimination' over $\mathrm{Mod}(T)$: every $\Diamond\phi \in L^{\Diamond}$ is equivalent in $\mathrm{Mod}(T)$ to ϕ. When T is $\aleph_0$-categorical, this can be seen as a special case of our Theorem 3.56, using the Fundamental Theorem of Bisimulations (Theorem 3.21). Indeed, if M is a world in a potentialist system which sees only the identity embedding, then for every $\phi \in L^{\Diamond}_{\infty}$ and $\overline{a}$ in M we have that $\mathcal{M}, M \vDash \phi[\overline{a}]$ if and only if $\mathcal{M}, M \vDash \Diamond\phi[\overline{a}]$.

Proof of Theorem 3.56. Define the system $\mathcal{N}$ as follows. Start with all the finite worlds of $\mathcal{M}$ together with the embeddings between them. Let Q be a countable model of T. Add Q to $\mathcal{N}$, together with the identity embedding. Furthermore, for every finite world M in $\mathcal{M}$ and embedding $\pi\colon M \to M'$ into an infinite world, by the Downward Lowenheim-Skolem Theorem we can realise π as an embedding $\pi^*\colon M \to Q$. Add π^* into $\mathcal{N}$. Note that multiple embeddings may result in the same starred versions which are added to $\mathcal{N}$.

Now define the relation $\sim$ between $\mathcal{M}$ and $\mathcal{N}$ from two cases (c.f. the proof of Theorem 3.53).

- **Case 1:** M **is finite**. Relate $(M, \overline{a})$ and $(N, \overline{b})$ if and only if they are equal.
- **Case 2:** M **is infinite**. Relate $(M, \overline{a})$ and $(N, \overline{b})$ if and only if $N = Q$ and $\overline{a}$ and $\overline{b}$ agree on all atomic formulas: $\Delta_0^M(\overline{a}) = \Delta_0^N(\overline{b})$.

It is not hard to see using the Downward Lowenheim-Skolem Theorem that this is a bitotal relation. Let us verify that it is a bisimulation. Assume that $(M, \overline{a}) \sim (N, \overline{b})$.

(B1) The first condition follows by definition.

(B2) When M is finite this condition is immediate, so assume that M is infinite. By the Ryll-Nardzewski Theorem applied to $\mathrm{Th}(M)$, the type of every tuple is principal. Moreover, since T has quantifier elimination, the complete types of $\overline{a}$ and $\overline{b}$ are the same. Hence any extension of one can be mirrored in the other.

(B3) There are several cases to consider. If M is infinite, then since any embedding preserves atomic types, any modal extension on one side can be mirrored on the other. So we can assume that M is finite (and hence $N = M$). Now any embedding from a finite world to a finite world on on side can be identically mirrored on the other, since $\mathcal{N}$ includes all embeddings from $\mathcal{M}$ between finite worlds. So we are left with two cases. (i) If we have $\pi\colon M \to M'$ in $\mathcal{M}$ with M' infinite, then as above we can realise this as $\pi^*\colon M \to Q$. Note that the atomic types of $\pi(\overline{a})$ and $\pi^*(\overline{b})$ are the same. (ii) If we have $\pi^*\colon N \to Q$ in $\mathcal{N}$, then by definition there is $\pi\colon M \to M'$ in $\mathcal{M}$ from which π^* was obtained. Again the atomic types of $\pi(\overline{a})$ and $\pi^*(\overline{b})$ are the same. □

Corollary 3.58. *Let T be an $\aleph_0$-categorical first-order L-theory which has quantifier elimination and no finite models. Then between any two potentialist systems which consist of models of T there is a bitotal bisimulation.*

A partial converse can be obtained to Theorem 3.56. In order to make it work, we need to make sure that $\mathcal{M}$ has 'enough structure' from $\mathrm{Mod}(T)$.

Theorem 3.59. *Let $\mathcal{M}$ be a potentialist system and let $T = \mathrm{Th}(\mathcal{M})$. Assume that $\mathcal{M}$ has a isomorphic copy of every countable model of T, together with all embeddings between these isomorphic copies. Assume further that $\mathcal{M}$ is bitotally bisimilar to a system $\mathcal{N}$ which contains exactly one infinite world which has no self-embeddings except the identity. Then T is $\aleph_0$-categorical and has quantifier elimination.*

Note that we don't need the infinite world in $\mathcal{N}$ to be countable.

Proof. Let Q be the unique infinite world in $\mathcal{N}$, and let $\sim$ be the bisimulation between $\mathcal{M}$ and $\mathcal{N}$. Note that if M in $\mathcal{M}$ is infinite, then every pair $(M, \overline{a})$ is related to a pair with Q.

First, let us see that T is $\aleph_0$-categorical. Take X, Y any countable models of T. We can assume without loss of generality that X and Y are in $\mathcal{M}$. We will build up an isomorphism $X \cong Y$ using a back-and-forth argument. First, using that $\sim$ is bitotal, for any $x_0 \in X$ the pair $(X, (x_0))$ is related to some pair $(Q, (q_0))$, which must in turn be related some pair $(Y, (y_0))$. Going the other direction, for any $y_1 \in Y$, the pair $(Y, (y_0, y_1))$ is related to some $(Q, (q_0, q_1))$, which is related to some $(X, (x_0, x_1))$. Proceeding in this way, using enumerations of X and Y, we obtain a bijection $X \to Y$. Using condition (B1) on bisimulations, each finite correspondence is a partial isomorphism. Hence the whole correspondence $X \to Y$ is an isomorphism.

Second, let us check that T admits quantifier elimination. First note that since T is $\aleph_0$-categorical, it is complete. Let r be any atomic type over T, and let p and q be extensions of r to a complete type. The goal is to show that $p = q$. Let $W \in \mathcal{M}$ be a countable model with $\overline{w}$ in W realising r. Since $\sim$ is bitotal, there is a pair $(Q, \overline{q})$ related to $(W, \overline{w})$. By a compactness argument, using that T is complete, there are countable models X and Y of T and embeddings $\pi\colon W \to X$ and $\rho\colon W \to Y$ such that $\pi(\overline{w})$ realises p and $\rho(\overline{w})$ realises q. We can assume that X, Y, π, ρ are in $\mathcal{M}$. Now, since in $\mathcal{N}$ the only embedding from Q is the identity, we have that:

$$(X, \rho(\overline{w})) \sim (Q, \overline{q}) \sim (Y, \rho(\overline{w}))$$

By the Fundamental Theorem of Bisimulations (Theorem 3.21), in particular:

$$X, \rho(\overline{w}) \equiv Y, \rho(\overline{w})$$

Hence $p = q$ as required. □

Let us see that it is necessary to assume that $\mathcal{M}$ has a isomorphic copy of every countable model of T. Let T be the theory of non-zero $\mathbb{Q}$-vector spaces, and let $\mathcal{M}$ be the subsystem of $\mathrm{Mod}(T)$ consisting of the uncountable worlds. Then $\mathcal{M}$ is bitotally bisimilar with the system consisting just of the $\aleph_0$-dimensional $\mathbb{Q}$-vector space with the identity embedding. However $T = \mathrm{Th}(\mathcal{M})$ is not $\aleph_0$-categorical.

I conclude this section with an example of how things can go wrong for Theorem 3.56 when we don't require that T admit quantifier elimination. The point is that once we are allowed to play with slightly more intricate first-order models, we can start arranging them to produce a rich modal structure.

Let $L = \{\approx\}$, and let T specify that $\approx$ is an equivalence relation with infinitely many size-1 and size-2 equivalence classes, and no other classes. Then T is $\aleph_0$-categorical, but does not have quantifier elimination. However every one-variable formula is equivalent to either an existential or a universal formula.

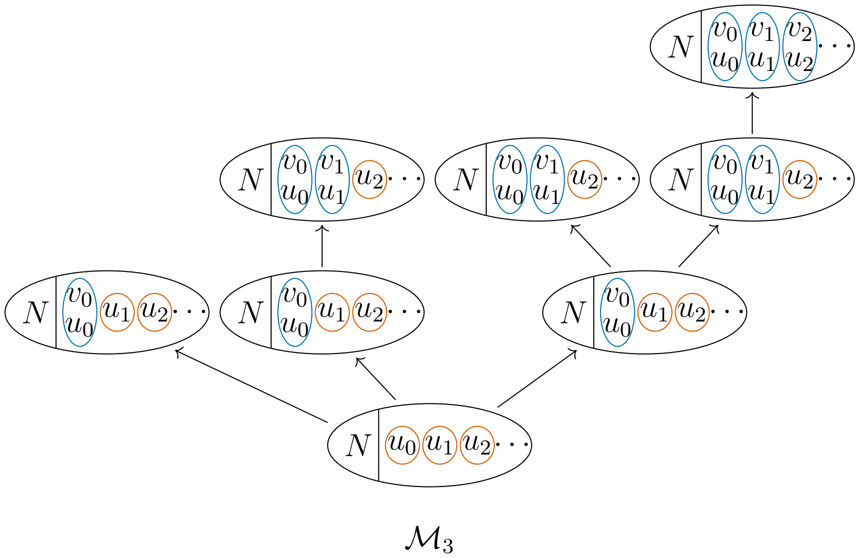

Figure 3.4: The potentialist system $\mathcal{M}_3$

We construct for every ordinal α a pointed system $(\mathcal{M}_\alpha, M_\alpha)$ of countable models of T. See Figure 3.4 for a representation of the system $\mathcal{M}_3$. The underlying modal structure of $\mathcal{M}_\alpha$ is the result of viewing α as a well-founded tree X_α. (That is, X_α consists of a root sitting below disjoint copies of each X_β for $\beta < \alpha$.) We place countable models of T along this structure as follows. First let N be any countable model. At the root of X_α we place a model consisting of N plus new elements $u_0, u_1, \ldots$, each in a singleton equivalence class. On the next level up we place disjoint copies of the same model with a new element v_0, which appears in an equivalence class with u_0. The model at the root embeds into those above it via the inclusion embedding. The construction proceeds upwards in this fashion. Let M_α be the model appearing at the root in this system.

Theorem 3.60. *Each $(\mathcal{M}_\alpha, M_\alpha)$ satisfies a different $L^\Diamond_\infty$-theory. Hence the system consisting of the disjoint union of each $\mathcal{M}_\alpha$ cannot be bitotally bisimilar with a set-sized system.*

Proof. The idea is that each pair (u_m, v_m) acts as a 'button'. The lowest model has all $\aleph_0$-many buttons unpushed. Each time we move up a level, we push an extra button (i.e. add the element v_m to make a 2-element equivalence class). Of course, each model is isomorphic, so the first-order structure cannot tell us how many buttons have been pushed, or how far up the tree we are. However, using the modal language, we can indeed talk about pushing a button, to get to a higher level.

Let:

$$p(x) := \neg\exists y(x \neq y \wedge x \approx y)$$
$$q(x) := \exists y(x \neq y \wedge x \approx y)$$

Define for each α the sentence θ_α by induction:

$$\theta_\alpha := \exists x\left(p(x) \wedge \bigwedge_{\beta<\alpha} \Diamond(q(x) \wedge \theta_\beta)\right) \wedge \forall x\left(p(x) \to \Box\left(q(x) \to \bigvee_{\beta<\alpha} \theta_\beta\right)\right)$$

In words: "there is a button which is pushed in future models satisfying θ_β for all $\beta < \alpha$, and whenever a button gets pushed in the future, that model satisfies θ_β for some $\beta < \alpha$". By induction we see that:

$$\mathcal{M}_\beta, M_\beta \vDash \theta_\alpha \quad \Leftrightarrow \quad \beta = \alpha$$

Therefore, each $(\mathcal{M}_\alpha, M_\alpha)$ satisfies a different $L^{\Diamond}_{\infty}$-theory.

Now, let $\mathcal{M}$ be the disjoint union of every $\mathcal{M}_\alpha$. If this were bitotally bisimilar with a set-sized system $\mathcal{N}$, then class-many of the M_α's would be related to the same world in N. But this is impossible by the Fundamental Theorem of Bisimulations (Theorem 3.21). □

3.7 Open questions

The first part of this chapter presented a number of foundational results concerning bisimulations of potentialist systems. The second part considered Question 3.42: when is a potentialist system $\mathcal{M}$ bisimilar with a set-sized one? In Section 3.5, I showed that if T is κ-rich for some κ then $\mathrm{Mod}(T)$ is bisimilar with a set-sized system. Does the converse hold? An iteration argument shows that if $\mathrm{Mod}(T)$ is bisimilar with a set-sized system, it is necessarily bisimilar with some $\mathrm{Mod}_\lambda(T)$. Will it further be bisimilar with some $\mathrm{Mod}_{\kappa^+}(T)$?

Section 3.6 seeks to give answers to Question 3.42 based solely on $\mathrm{Th}(\mathcal{M})$. This is a rather crude measure, and takes no account of the modal structure of $\mathcal{M}$. Theorem 3.56 shows that if $\mathrm{Th}(\mathcal{M})$ is $\aleph_0$-categorical and admits quantifier elimination, then $\mathcal{M}$ is bisimilar with a set-sized system. Is this the best result possible? If instead of $\mathrm{Th}(\mathcal{M})$ we consider the set of (finitary) first-order modal sentences satisfied by $\mathcal{M}$, can we obtain a more general characterisation?

Bibliography

[Ada22a] Sam Adam-Day. *Bisimulations of potentialist systems*. Preprint submitted to the *Journal of Symbolic Logic*. 2022. arXiv: 2206.10359 [math.LO].

[Ada22b] Sam Adam-Day. 'On the continuous gradability of the cut-point orders of R-trees'. In: *Topology and its Applications* 306 (2022), p. 107937. ISSN: 0166-8641. DOI: 10.1016/j.topol.2021.107937.

[Ada23] Sam Adam-Day. *Uniform, rigid branchwise-real trees*. To appear in the *Israel Journal of Mathematics*. 2023. arXiv: 2206.15344 [math.LO].

[AS85] Uri Abraham and Saharon Shelah. 'Isomorphism types of Aronszajn trees'. In: *Israel Journal of Mathematics* 50 (1985), pp. 75–113. DOI: 10.1007/BF02761119.

[Avr79] Uri Avraham. 'Construction of a rigid Aronszajn tree'. In: *Proceedings of the American Mathematical Society* 77.1 (1979). DOI: 10.1090/S0002-9939-1979-0539646-8.

[Bau70] James E. Baumgartner. 'Results and Independence Proofs in Combinatorial Set Theory'. PhD thesis. University of California, 1970. ISBN: 9798659828038.

[Ben10] Johan van Benthem. *Modal logic for open minds*. CSLI lecture notes 199. Stanford, California: CSLI, 2010. ISBN: 9781575865980.

[Ben14] Johan van Benthem. *Logic in Games*. The MIT Press, 2014. ISBN: 9780262320290. DOI: 10.7551/mitpress/9674.001.0001.

[Ber19] Valerii Nikolaevich Berestovskii. 'On Urysohn's R-Tree'. In: *Siberian Mathematical Journal* 60.1 (2019), pp. 10–19. DOI: 10.1134/S0037446619010026.

[Bes01] Mladen Bestvina. 'Real trees in topology, geometry, and group theory'. In: *Handbook of Geometric Topology*. Ed. by R.J. Daverman and R.B. Sher. Amsterdam: North-Holland, 2001, pp. 55–91. ISBN: 978-0-444-82432-5. DOI: 10.1016/B978-044482432-5/50003-2. arXiv: math/9712210 [math.GT].

[Bow99] B.H. Bowditch. *Treelike structures arising from continua and convergence groups*. Vol. 139. Memoirs of the American Mathematical Society 662. E-print available: http://homepages.warwick.ac.uk/~masgak/papers/bhb-treelike.pdf. 1999. ISBN: 978-0-8218-1003-3. DOI: 10.1090/memo/0662.

[BRV01] Patrick Blackburn, Maarten de Rijke and Yde Venema. *Modal logic*. Cambridge tracts in theoretical computer science 53. Cambridge: Cambridge University Press, 2001. ISBN: 9780521802000. DOI: 10.1017/CBO9781107050884.

[Can95] Georg Cantor. 'Beiträge zur Begründung der transfiniten Mengenlehre'. In: *Mathematische Annalen* 46 (1895), pp. 481–512. DOI: 10.1007/BF02124929.

[Chi01] Ian Chiswell. *Introduction to Λ-trees*. Singapore, 2001. ISBN: 9789810243869. DOI: 10.1142/4495.

[CZ97] Alexander Chagrov and Michael Zakharyaschev. *Modal logic*. Oxford logic guides 35. Oxford: Clarendon Press, 1997. ISBN: 9780198537793.

[DP01] Anna Dyubina and Iosif Polterovich. 'Explicit Constructions of Universal R-Trees and Asymptotic Geometry of Hyperbolic Spaces'. In: *Bulletin of the London Mathematical Society* 33.6 (Nov. 2001), pp. 727–734. DOI: 10.1112/S002460930100844X. arXiv: math/9904133 [math.DG].

[ET43] P. Erdös and A. Tarski. 'On Families of Mutually Exclusive Sets'. In: *Annals of Mathematics* 44.2 (1943), pp. 315–329. ISSN: 0003486X. DOI: 10.2307/1968767.

[Fab15] Paul Fabel. 'A Topological Characterization of the Underlying Spaces of Complete R-Trees'. In: *Michigan Mathematical Journal* 64 (2015). E-print available: https://projecteuclid.org/download/pdf_1/euclid.mmj/1447878035, pp. 881–887. DOI: 10.1307/mmj/1447878035.

[FH09] Gunter Fuchs and Joel David Hamkins. 'Degrees of Rigidity for Souslin Trees'. In: *The Journal of Symbolic Logic* 74.2 (2009), pp. 423–454. ISSN: 00224812, 19435886. DOI: 10.2178/jsl/1243948321. arXiv: math/0602482 [math.LO].

[FJ04] Charles Favre and Mattias Jonsson. *The Valuative Tree*. Lecture notes in mathematics 1853. Berlin; London: Springer, 2004. ISBN: 9783540229841. DOI: 10.1007/b100262. arXiv: math/0210265 [math.AC].

[FM98] Melvin Fitting and Richard L Mendelsohn. *First-order modal logic*. Synthese library 277. Dordrecht; London: Kluwer Academic, 1998. ISBN: 978-0-7923-5334-8. DOI: 10.1007/978-94-011-5292-1.

[GS64] Haim Gaifman and E. Specker. 'Isomorphism Types of Trees'. In: *Proceedings of The American Mathematical Society* 15 (Feb. 1964), pp. 1–7. DOI: 10.1090/S0002-9939-1964-0168484-2.

[GSS09] Dov M Gabbay, Dimitrij Skvortsov and Valentin Shehtman. *Quantification in Nonclassical Logic*. Vol. 1. Studies in logic and the foundations of mathematics. Oxford: Elsevier Science, 2009. ISBN: 9780444520128.

[Hal17] Lorenz J. Halbeisen. *Combinatorial Set Theory. With a Gentle Introduction to Forcing*. 2nd ed. Springer Monographs in Mathematics. Springer Cham, 2017. ISBN: 978-3-319-60230-1. DOI: 10.1007/978-3-319-60231-8.

[Ham03] Joel David Hamkins. 'A simple maximality principle'. In: *Journal of Symbolic Logic* 68.2 (2003), pp. 527–550. DOI: 10.2178/jsl/1052669062. arXiv: math/0009240 [math.LO].

[Ham18] Joel David Hamkins. *The modal logic of arithmetic potentialism and the universal algorithm*. 2018. arXiv: 1801.04599 [math.LO].

[HC96] G. E. Hughes and M. J. Cresswell. *A new introduction to modal logic*. London: Routledge, 1996. ISBN: 9780415125994.

[HL08] Joel David Hamkins and Benedikt Löwe. 'The modal logic of forcing'. In: *Transactions of the American Mathematical Society* 360.4 (2008), pp. 1793–1817. ISSN: 0002-9947. DOI: 10.1090/S0002-9947-07-04297-3. arXiv: math/0509616 [math.LO].

[HL13] Joel David Hamkins and Benedikt Löwe. 'Moving up and down in the generic multiverse'. In: *Logic and its Applications. ICLA 2013.* Ed. by Kamal Lodaya. Vol. 7750. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013, pp. 139–147. DOI: 10.1007/978-3-642-36039-8_13. arXiv: 1208.5061 [math.LO].

[HL19] Joel David Hamkins and Øystein Linnebo. 'The modal logic of set-theoretic potentialism and the potentialist maximality principles'. In: *Review of Symbolic Logic* (Oct. 2019), pp. 1–36. DOI: 10.1017/S1755020318000242. arXiv: 1708.01644 [math.LO].

[HLL15] Joel David Hamkins, George Leibman and Benedikt Löwe. 'Structural connections between a forcing class and its modal logic'. In: *Israel Journal of Mathematics* 207.2 (2015), pp. 617–651. DOI: 10.1007/s11856-015-1185-5. arXiv: 1207.5841 [math.LO].

[Hod93] Wilfrid Hodges. *Model theory*. Encyclopedia of mathematics and its applications 42. Cambridge: Cambridge University Press, 1993. ISBN: 9780521304429. DOI: 10.1017/CBO9780511551574.

[HW18] Joel David Hamkins and W. Hugh Woodin. *The universal finite set*. 2018. arXiv: 1711.07952 [math.LO].

[HW20] Joel David Hamkins and Wojciech Aleksander Wołoszyn. *Modal model theory*. 2020. arXiv: 2009.09394 [math.LO].

[HW21] Joel David Hamkins and Kameryn J. Williams. 'The Σ1-definable universal finite sequence'. In: *The Journal of Symbolic Logic* (2021), pp. 1–19. DOI: 10.1017/jsl.2020.59. arXiv: 1909.09100 [math.LO].

[Jec03] Thomas Jech. *Set Theory*. The Third Millennium Edition. Springer Monographs in Mathematics. Springer-Verlag, 2003. ISBN: 978-3-540-44085-7. DOI: 10.1007/3-540-44761-X.

[Jec72] Thomas J. Jech. 'Automorphisms of ω1-Trees'. In: *Transactions of the American Mathematical Society* 173 (1972), pp. 57–70. DOI: 10.1090/S0002-9947-1972-0347605-1.

[Jen69] Ronald Björn Jensen. 'Automorphism properties of Souslin continua'. In: *Notices of the American Mathematical Society* 16 (1969), p. 576.

[Jon65] P. Burton Jones. 'Remarks on the normal Moore space metrization problem'. In: *Topology Seminar Wisconsin*. Vol. 60. Princeton: Princeton University Press, 1965, pp. 115–120. ISBN: 9781400882076. DOI: 10.1515/9781400882076-015.

[Kun13] Kenneth Kunen. *Set Theory*. Revised edition. Studies in logic 34. London: College Publications, 2013. ISBN: 9781848900509.

[Lin13] Øystein Linnebo. 'The potential hierarchy of sets'. In: *The Review of Symbolic Logic* 6.2 (2013), pp. 205–228. DOI: 10.1017/S1755020313000014.

[MNO92] John C. Mayer, Jacek Nikiel and Lex G. Oversteegen. 'Universal Spaces for R-Trees'. In: *Transactions of the American Mathematical Society* 334.1 (1992), pp. 411–432. ISSN: 00029947. DOI: 10.2307/2153989.

[Mor92] John W Morgan. 'Λ-Trees and Their Applications'. In: *Bulletin of the American Mathematical Society* 26 (1992), pp. 87–113. DOI: 10.1090/S0273-0979-1992-00237-9. arXiv: math/9201265 [math.GR].

[Nik89] Jacek Nikiel. *Topologies on pseudo-trees and applications*. Vol. 416. Memoirs of the American Mathematical Society. American Mathematical Society, 1989. ISBN: 9780821824795. DOI: 10.1090/memo/0416.

[Rud76] Walter Rudin. *Principles of Mathematical Analysis*. 3rd Edition. International series in pure and applied mathematics. New York: McGraw-Hill, 1976. ISBN: 9780070542358.

[Sha87] Peter B. Shalen. 'Dendrology of Groups: An Introduction'. In: *Essays in Group Theory*. Ed. by S. M. Gersten. New York, NY: Springer New York, 1987, pp. 265–319. ISBN: 978-1-4613-9586-7. DOI: 10.1007/978-1-4613-9586-7_4.

[Sha91] Peter B. Shalen. 'Dendrology and its applications'. In: *Proceedings of the Group Theory from a Geometrical Viewpoint, ICTP*. Ed. by E Ghys, A Haefliger and A Verjovsky. E-print available: http://homepages.math.uic.edu/~shalen/shebang.pdf. Trieste, Italy: World Scientific, 1991. DOI: 10.1142/1235.

[SS20] Denis I. Saveliev and Ilya B. Shapirovsky. 'On modal logics of model-theoretic relations'. In: *Studia Logica* 108 (2020), pp. 989–1017. DOI: 10.1007/s11225-019-09885-y. arXiv: 1804.09810 [math.LO].

[Ury27] Paul Urysohn. 'Beispiel eines nirgends separablen metrischen Raumes'. German. In: *Fundamenta Mathematicae* 9 (1927), pp. 119–121. DOI: 10.4064/fm-9-1-119-121.

[Woł22] Wojciech Aleksander Wołoszyn. 'Modal theory of the category of sets'. Preprint in preparation. 2022.