Self-embedding monoids and the existential Thomas conjecture

András Pongrácz
University of Debrecen, Institute of Mathematics
4032 Debrecen, Egyetem tér 1, Hungary
ORCID: 0000-0002-2771-8974

14 November 2026

hexagon:2611.00444v1

Abstract

We give a countable homogeneous finitely bounded Ramsey structure in a finite relational language with finitely many reducts up to first-order interdefinability but infinitely many self-embedding monoids of reducts. The reducts producing these monoids are first-order interdefinable, while their monoids form a strictly increasing chain. This refutes the existential strengthening of Thomas' conjecture even under the Ramsey hypothesis. The example is an ordered expansion of an equivalence-relation structure. Before the order is added, that structure has exactly five reducts up to first-order interdefinability.

We also give a revised proof of the classification of self-embedding monoids for reducts of the countable existentially closed binary branching semilinear order. Exactly five monoids occur, refining its three-class first-order classification. The proof transfers the known classification for the homogeneous binary branching C-relation from antichains to arbitrary finite configurations. The additional step is a connectivity argument on finite labelled trees, using both preservation and reflection of the reduct's basic relations.

2020 Mathematics Subject Classification. Primary 03C15; Secondary 20B27, 03C40, 20M20.
Keywords. Semilinear order, reduct, self-embedding monoid, existential definability, binary tree, C-relation.

1 Introduction

A *reduct* of a relational structure Δ is a structure on the same domain whose basic relations are first-order definable in Δ without parameters. For a countable ω-categorical structure, classifying reducts up to first-order interdefinability is equivalent to classifying permutation groups containing $\mathrm{Aut}(\Delta)$ that are closed in the pointwise-convergence topology on the symmetric group of the domain. Thomas conjectured that this interval is finite whenever Δ is homogeneous in a finite relational language [39, 40]. Classical classifications include the rational order [18], the random graph and random hypergraphs [39, 40], and expansions of the rational order by constants [27]. Further classifications cover the random partial order [35], the random ordered graph [15], the random permutation [32], the generic digraph [3], the Henson digraphs [4], and Henson graphs expanded by a constant [36]. Bodirsky's monograph [6] provides general background on definability, oligomorphic groups and their transformation monoids.

Self-embeddings retain more information than automorphisms. An embedding preserves each basic relation and its complement, and an ω-categorical structure's self-embedding monoid determines all of its existentially definable relations.

Recall that a structure is *model-complete* if every first-order formula is equivalent, modulo its theory, to an existential formula. For a countable ω-categorical structure Γ with domain D, the criterion due to Bodirsky and Pinsker, recorded with that attribution by Bodirsky and Junker [11, Corollary 2.3], states that Γ is model-complete exactly when $\mathrm{Emb}(\Gamma) = \overline{\mathrm{Aut}(\Gamma)}$. Here the closure is taken in D^D with the topology of pointwise convergence, with D discrete. The equality means that every self-embedding agrees with an automorphism on each finite subset of D. A *core* is a structure whose endomorphisms are embeddings. Consequently, writing $\mathrm{End}(\Gamma)$ for the endomorphism monoid, Γ is a model-complete core exactly when $\mathrm{End}(\Gamma) = \overline{\mathrm{Aut}(\Gamma)}$; see also [5, Theorem 18]. Thus for model-complete reducts the automorphism group already determines the self-embedding monoid; for example, all reducts of the random graph are model-complete [14].

The existential strengthening of Thomas' conjecture asks whether a countable homogeneous structure in a finite relational language has only finitely many self-embedding monoids of reducts. The author's earlier manuscript formulated this assertion *with the Ramsey hypothesis* [37, Conjecture 4.5]; the broader formulation is discussed in [38]. Its definability formulation requires existential definitions of both basic relations and their complements; Section 3 explains this convention.

Our counterexample refutes both formulations. Recall that a structure is *Ramsey* if, for any finite substructures A, B and integer $r \geq 2$, some finite substructure C has the property that every r-colouring of the copies of A in C is constant on those contained in some copy of B. It is *finitely bounded* if its age is specified by a finite list of forbidden induced substructures.

Theorem 1.1. *There exists a countably infinite homogeneous, finitely bounded Ramsey structure in a finite relational language of maximum arity two which has finitely many reducts up to first-order interdefinability, but admits a sequence of quantifier-free reducts whose self-embedding monoids form a strictly increasing infinite chain. These reducts are all first-order interdefinable with one another, and each has a single basic relation.*

The construction starts with an equivalence relation having infinitely many singleton classes and infinitely many infinite classes. Adding a unary predicate for the set of nonsingleton-class points makes the structure homogeneous. Its first-order reducts form a five-element lattice, which we determine explicitly; their finiteness also follows from Bodor's general theorem [16]. A suitable convex order gives the Ramsey expansion in Theorem 1.1. The infinite chain is constructed in Theorem 7.1, and Proposition 7.3 verifies homogeneity, finite boundedness and the Ramsey property. For this ordered expansion, Appendix B proves first-order finiteness, without enumerating its reducts, using Cameron's classification [18] and Bodor's bounds for finite covers [17].

The other part of the paper concerns the binary branching semilinear order. Its first-order reducts were classified by Bodirsky, Bradley-Williams, Pinsker and Pongrácz [8]; their analysis also identified an injection into an antichain that preserves and reflects a derived rooted-tree relation, but does not preserve the original order. Thus an automorphism-group classification does not settle the embedding-monoid problem. The five-monoid classification was stated in the author's unpublished 2020 manuscript [37, Theorem 2.7]. Here we give a revised and completed proof, making the reconstruction from homogeneous leaf structures explicit. The five monoids form the lattice in Figure 4. Three come from the original order, its betweenness relation, and equality. The remaining two preserve the rooted-triple relation R and the unrooted-quartet relation P obtained by replacing each marked point by a leaf. These relations will be defined explicitly in Section 2.

The structure and automorphism groups of semilinear orders were studied by Droste [20, 21] and by Droste, Holland and Macpherson [22, 23, 24]. Semilinear orders and their associated betweenness relations also occur among the structures in Adeleke and Macpherson's

classification of infinite primitive Jordan permutation groups [1]; Macpherson [33] surveys this connection with model theory. For an earlier study of structures described by relations on tree leaves, see Cameron [19]. Adeleke and Neumann [2] give a systematic treatment of semilinear orders, betweenness relations, and the ternary branching and quaternary separation relations known as C- and D-relations. Section 2 specifies the conventions for the binary branching relations used here.

The proof uses the published classification of reducts of the homogeneous binary branching C-relation by Bodirsky, Jonsson and Pham [9]. Passing to a maximal antichain gives access to that theorem, but it does not by itself reconstruct the monoid on the whole semilinear order. The missing step concerns maps that preserve every rooted triple on antichains but change a triple containing comparable points. We show that one such map forces the basic relations of the reduct to depend only on equality. The argument proceeds through finite labelled trees: a local change supplied by the map can be inserted into larger configurations, and induction connects all trees on a fixed set of labels. The corresponding unrooted argument is reduced to this rooted argument by fixing an auxiliary leaf.

A map is called *flat* if it is injective and its image is an antichain. The initial reduction to such a map is inherited from [8]; the leaf classification is inherited from [9]. The finite-tree and rerooting techniques have predecessors in [9, Section 4.2] and [8, Lemma 4.27]; the two-sided invariant-relation argument below is what allows us to use them for self-embedding monoids. The counterexample and its first-order reduct analysis are additional to the author's unpublished 2020 manuscript. For general results on homogeneous images of canonical maps and transfer of the Ramsey property, see Mottet and Pinsker [34]. Structural Ramsey theory enters through the cited classification and flat-map results. Appendix A supplies a direct proof of the specific Ramsey input used there, avoiding an invalid general transfer assertion in the original citation chain. The reconstruction arguments require no additional canonisation theorem.

These classifications also connect with constraint satisfaction through primitive-positive definability. A *primitive-positive* (pp) formula is an existentially quantified conjunction of atomic formulas. For a structure Γ with a finite relational signature, CSP(Γ) asks whether an input pp sentence holds in Γ. For ω-categorical Γ, a nonempty relation is pp-definable precisely when it is preserved by every *polymorphism* of Γ, that is, every operation of finite positive arity preserving the basic relations coordinatewise, by the preservation theorem of Bodirsky and Nešetřil [13]; see also [6]. Consequently pp-interdefinable finite-signature structures have polynomial-time equivalent CSPs. The relevant invariant is therefore the collection of all polymorphisms, rather than the self-embedding monoid alone. For the random partial order, Kompatscher and Pham build on the first-order classification [35], analyse endomorphisms and polymorphisms, and obtain a polynomial-time versus NP-complete dichotomy for every finite-signature reduct [29, Theorems 1 and 4]. This establishes the complexity classification without requiring an enumeration of all pp-interdefinability classes. Related dichotomies hold for temporal constraint satisfaction, that is, reducts of the rational order [12], and for phylogeny constraint satisfaction, that is, reducts of the homogeneous binary branching C-relation [10]. In the latter case the existential classification [9], also used here, supplies a structural input to the complexity analysis.

2 Structures, tree relations and the classification

2.1 The semilinear order and its finite diagrams

A partial order is *semilinear* if any two elements have a common upper bound and the upper bounds of each element form a chain. We work with the unique countable existentially closed

semilinear order $(S_2; \leq)$. We use the description and homogeneity result of [8, Proposition 3.4 and Theorem 3.5]. In particular, the order is dense, unbounded, binary branching, nice, and without joins. Here $x \parallel y$ means that x and y are incomparable. Unboundedness means that for every x there are y, z with $y < x < z$. We shall use the following forms of the other properties:

1. *Density and branching below a point:* if $x < y$, some z satisfies $x < z < y$, and some $u < y$ is incomparable with x. The latter condition is the first binary-branching requirement.

2. *Niceness:* if $x \parallel y$, some $z > x$ is incomparable with y.

3. *Absence of joins:* if $x, y < z$ and $x \parallel y$, there is a common upper bound of x, y strictly below z.

4. *Binary branching on antichains:* among three incomparable points, exactly one has the following property: the other two have a common upper bound incomparable with that point. Existence is the second binary-branching requirement; uniqueness follows from semilinearity.

The structure is universal for countable semilinear orders.

For pairwise distinct x, y, z, put

$$C(x; y, z) \quad \Longleftrightarrow \quad x, y, z \text{ are pairwise incomparable and } \exists u\,(y < u \wedge z < u \wedge u \parallel x).$$

Set $C(x; y, z)$ to false on all other triples. Then $\Delta = (S_2; \leq, C)$ is homogeneous, ω-categorical, and has quantifier elimination. Its self-embeddings are exactly those of $(S_2; \leq)$, by [8, Proposition 3.4 and the proof of Proposition 4.17]. The latter assertion also follows from model-completeness of the existentially closed semilinear order.

A *rooted binary leaf tree* is a finite rooted tree in which each non-leaf vertex has two children; the one-leaf tree is allowed. We regard the root as lying above its descendants. For pairwise distinct leaves x, y, z, write $x \mid yz$ if the least common ancestor of y, z lies strictly below the least common ancestor of x, y, z; set $x \mid yz$ to false when any two entries coincide. Labels on leaves are part of the structure. In the relation $x \mid yz$, the leaf x is called the *outgroup*. For three distinct leaves exactly one of $x \mid yz$, $y \mid xz$, $z \mid xy$ holds.

We will also use *normalised finite tree diagrams*: finite rooted trees with at most two children per vertex, with labels at some vertices, and with every branching vertex unlabelled. Thus a labelled non-leaf vertex has just one child and lies on a nonbranching part of the diagram; it may also be the root. Unlabelled vertices with one child are allowed. Equip the full vertex set with the ancestor order, with ancestors above descendants, and interpret C by the preceding formula in that finite order. The induced relations on the labelled vertices are what the diagram prescribes. In particular, auxiliary unlabelled vertices may witness C. All tree diagrams below use this convention. Figure 1 illustrates the convention and the order and C-relations read from a diagram. The next observation spells out how such diagrams are realised; this will be needed when a small configuration is kept fixed.

Lemma 2.1 (Realisation of finite tree diagrams)**.** *A normalised finite tree diagram can be realised in S_2 with the prescribed order and prescribed C-relations on its marked vertices. More precisely, any prescribed embedding into Δ of the induced $(\leq, C)$-structure on a subset A of its marked vertices extends to such a realisation of all marked vertices. Every nonempty finite subset of Δ has such a diagram.*

Proof. First consider a given diagram. Embed its full ancestor order into S_2 recursively. Choose the root image arbitrarily. For one child choose a smaller point; for two children choose

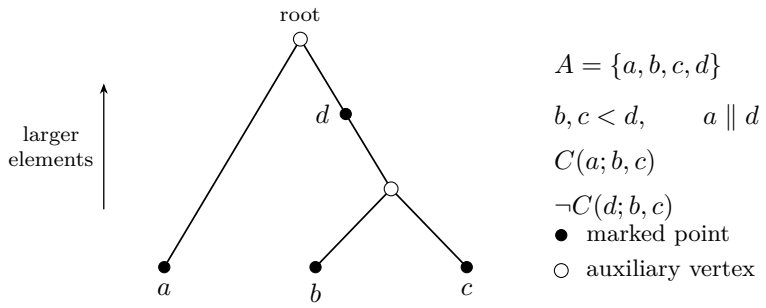


Figure 1: A normalised finite tree diagram on the marked set $A = \{a, b, c, d\}$. Filled vertices are marked, while open vertices are auxiliary and unlabelled. The interior marked point d has one child; both branching vertices are unmarked. The lower open vertex witnesses $C(a; b, c)$, whereas $C(d; b, c)$ fails because $b, c < d$. The auxiliary vertices specify upper bounds in a realisation, without requiring them to be joins in S_2.

two incomparable smaller points, and continue below them. Points in different child branches remain incomparable: otherwise a lower point would have two incomparable upper bounds, contrary to semilinearity. A common ancestor in the diagram witnesses each prescribed C-orientation. Exactly one orientation holds on an antichain triple in both structures, and C is false on all other triples. Thus the marked $(\leq, C)$-structure is preserved and reflected. For a prescribed embedding on A, the map from the realised copy of A to its prescribed copy is a finite $(\leq, C)$-isomorphism. Extend it to an automorphism of Δ by homogeneity and compose our realisation with that automorphism.

For completeness, we construct a normalised diagram of a nonempty finite set $A \subseteq S_2$. Consider the distinct nonempty sets

$$\mathcal{F}_A = \{A \cap {\downarrow}x : x \in S_2\} \setminus \{\varnothing\}, \qquad {\downarrow}x = \{y : y \leq x\}.$$

This is a finite laminar family: if two such sets meet at a, their defining upper points are comparable, so one set contains the other. It contains A, since A has a common upper bound. Its inclusion diagram is therefore a rooted tree. Mark the vertex $A \cap {\downarrow}a$ by a, for each $a \in A$; different points mark different vertices.

There are at most two children at each vertex. Indeed, suppose a cluster $X = A \cap {\downarrow}x$ has three maximal proper subclusters X_1, X_2, X_3, and choose $a_i \in X_i$. These points are pairwise incomparable. Binary branching gives an upper bound u of two of them incomparable with the third, say a_1. Since u and x have a common lower point, they are comparable; $x \leq u$ would contradict $u \parallel a_1$, so $u < x$. The cluster $A \cap {\downarrow}u$ is then a proper subcluster of X meeting two different maximal children. By laminarity it contains both of them, a contradiction.

A marked vertex has at most one child. For the vertex marked a, all points in $(A \cap {\downarrow}a) \setminus \{a\}$, if this set is nonempty, have a common strict upper bound $u < a$. To obtain it, start with an upper point of the first member below a, using density. On adding another member, if it is comparable with the current upper point, choose a point above their maximum and below a; if they are incomparable, use the no-joins property to find a common upper bound below a. Each step retains a strict upper bound for all members already treated, and the procedure ends because the set is finite. Their whole set is therefore a single cluster $A \cap {\downarrow}u$, which is the unique maximal proper subcluster of $A \cap {\downarrow}a$. Hence every branching vertex is unmarked, as required.

The order of marked vertices is the original order. For an antichain triple a, b, c, $C(a; b, c)$ holds precisely when a cluster contains b, c and excludes a. A witness to C gives such a cluster. Conversely, a point defining such a cluster is above b, c and incomparable with a,

since these three marked points form an antichain. This proves that the diagram has exactly the original C-relation. □

2.2 Leafification, rooted triples and unrooted quartets

Write $\mathrm{Dist}(x_1, \ldots, x_n)$ for pairwise distinctness. Define

$$R(x; y, z) \iff \mathrm{Dist}(x, y, z) \wedge \big[C(x; y, z) \vee (y < x \wedge z < x) \vee (x \parallel y \wedge x \parallel z \wedge (y < z \vee z < y))\big]. \tag{1}$$

This is the injective part of the relation of [8, Definition 4.22], with the coordinates cyclically permuted. The restriction to distinct triples is essential. Without it, the displayed order formula would also encode $y < x$ through the repeated triple (x, y, y), and could not be preserved by an injection into an antichain. We therefore make the restriction explicit throughout.

Lemma 2.2 (Leafification)**.** *For every nonempty finite $A \subseteq S_2$, the relation R on A is the rooted-triple relation of a labelled rooted binary leaf tree. Every finite labelled rooted binary leaf tree occurs in this way on an antichain. If L is a maximal antichain in S_2, there exists an embedding*

$$f : (S_2; R) \longrightarrow (L; C).$$

It may be chosen to fix any prescribed finite subset of L.

Proof. Use the diagram constructed in the proof of Lemma 2.1. Its leaves are marked: every minimal nonempty cluster contains the cluster of one of its marked points and hence equals it. Suppress unnecessary unlabelled vertices with one child. If a marked point is not already a leaf, attach a new leaf at that point, transfer its label to the new leaf, and retain the old point as an unlabelled vertex. Normalisation ensures that a marked point had at most one downward branch before this operation. The resulting tree is binary. Figure 2 shows this operation on the diagram from Figure 1.

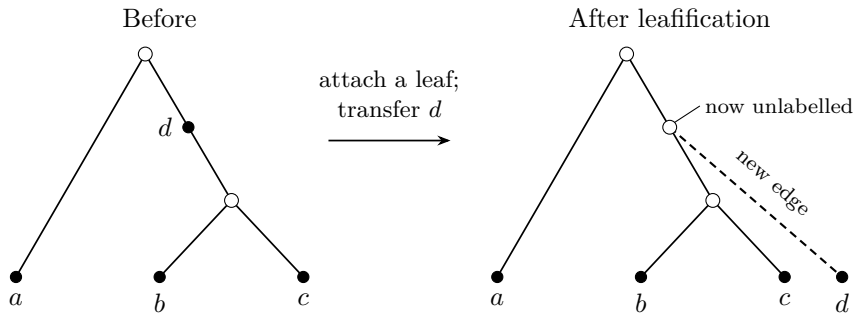


Figure 2: Leafification of Figure 1. A new leaf receives the label d; its former vertex becomes unlabelled. The other labels and all old edges stay in place, and the dashed edge is newly added. The leaf tree represents R on the original marked set: $b, c < d$ gives $R(d; b, c)$, now read as $d \mid bc$, while $a \mid bc$ is unchanged.

For three marked points that were incomparable, the operation retains their original C-orientation. If two are comparable and the third is incomparable with both, the third is the outgroup. If one point lies above the other two, that upper point is the outgroup, whether the other two are comparable or not. These exhaust the possible order patterns and are exactly the clauses of (1). This proves the first assertion. The converse follows by realising the leaves of a binary tree as an antichain using Lemma 2.1.

By [8, Section 2], a maximal antichain with its induced C-relation is the homogeneous universal binary branching leaf structure of [9], in the injective convention. Enumerate S_2. The first assertion says that every finite initial segment, with its R-relation, embeds into L. Such embeddings can be chosen compatibly: embed the next finite segment into L, then use homogeneity of L to align its previous segment with the embedding already chosen. Their union is the required f. To fix a finite subset of L, start this construction with its identity map. □

We call an injection *flat* if its image is an antichain. Thus the map in Lemma 2.2, regarded as a map into S_2, is a flat self-embedding of $(S_2; R)$.

For four distinct points define

$$P(x, y; u, v) \iff (R(x; u, v) \wedge R(y; u, v)) \vee (R(u; x, y) \wedge R(v; x, y)). \tag{2}$$

On tuples with repeated entries set P to false. On a leaf structure, define Q by the same formula with C in place of R on four pairwise distinct leaves, and set Q to false on quadruples with repeated entries. Geometrically, forget the root, suppress the resulting vertex of degree two, and take the unrooted binary leaf tree. Then $Q(x, y; u, v)$ means that its quartet split is $xy \mid uv$: the paths joining x to y and u to v are disjoint. For any four distinct leaves exactly one of the three quartet splits holds. For any map f supplied by Lemma 2.2, the defining formulas give

$$P(a, b; c, d) \iff Q(f(a), f(b); f(c), f(d)) \qquad (a, b, c, d \in S_2).$$

We abbreviate this coordinatewise equivalence as $P = f^{-1}(Q)$. Figure 3 illustrates this construction for the leaf tree in Figure 2.

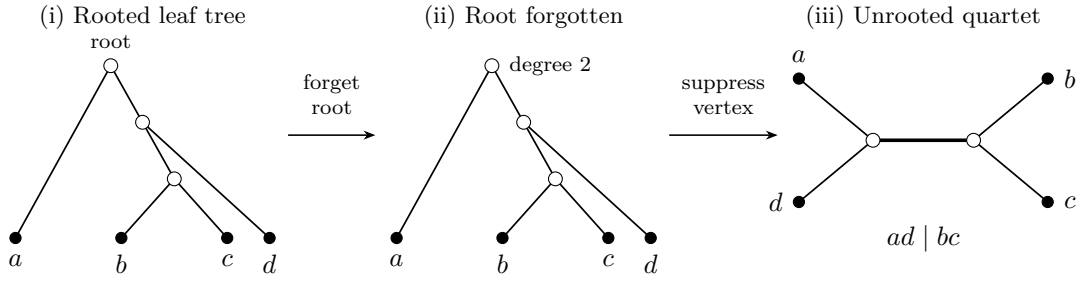


Figure 3: From the leaf tree in Figure 2 to an unrooted quartet. Forgetting the root leaves the top vertex with degree two. Suppressing it replaces its two incident edges by one edge and preserves all four leaves. In the final drawing, removing the thick central edge separates $\{a, d\}$ from $\{b, c\}$, giving $ad \mid bc$. Thus $Q(a, d; b, c)$ holds on the leaves, corresponding to $P(a, d; b, c)$ on the original marked set.

For comparison with [9, arXiv version, Sections 3.2–3.4], write $C^{\mathrm{std}}, Q^{\mathrm{std}}$ for that paper's leaf relations. They include repeated-coordinate cases, whereas our C, Q do not. The translations are

$$\begin{aligned} C^{\mathrm{std}}(x; y, z) &\iff C(x; y, z) \vee (x \neq y \wedge y = z), \\ Q^{\mathrm{std}}(x, y; u, v) &\iff Q(x, y; u, v) \\ &\quad \vee \big[(x = y \vee u = v) \wedge x \neq u \wedge x \neq v \wedge y \neq u \wedge y \neq v\big]. \end{aligned}$$

Conversely, our relations are the restrictions of $C^{\mathrm{std}}, Q^{\mathrm{std}}$ to pairwise distinct entries. These quantifier-free translations preserve the finite partial isomorphisms and the self-embeddings. Thus the cited homogeneity and embedding classifications apply to our conventions as well. The order formula (1) retains its explicit distinctness guard.

The semilinear betweenness relation is

$$B(x,y,z) \quad \Longleftrightarrow \quad (x<y<z)\vee(z<y<x)\vee(x<y\wedge y\parallel z)\vee(z<y\wedge y\parallel x).$$

In a normalised diagram it records that the marked point y lies on the open path from x to z in a finite tree diagram.

Write $\mathrm{Emb}(A,B)$ for the set of embeddings from a structure A into a structure B of the same signature, and $\mathrm{Emb}(\Theta)=\mathrm{Emb}(\Theta,\Theta)$ for the self-embedding monoid of Θ. Write $\mathrm{Inj}(D)$ for the monoid of all injections $D\to D$.

Theorem 2.3. *Let Γ be any first-order reduct of $\Delta=(S_2;\leq,C)$. Then its self-embedding monoid is exactly one of*

$$\mathrm{Emb}(S_2;\leq,C),\qquad \mathrm{Emb}(S_2;B),\qquad \mathrm{Emb}(S_2;R),\qquad \mathrm{Emb}(S_2;P),\qquad \mathrm{Inj}(S_2).$$

These five monoids are distinct, and their inclusions are precisely those in Figure 4.

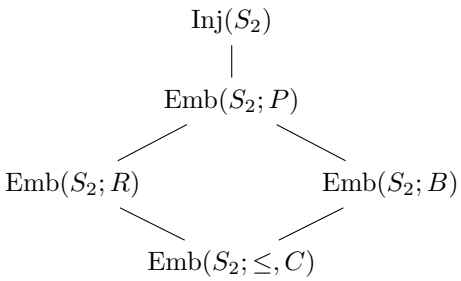


Figure 4: The five self-embedding monoids, ordered by inclusion. The two middle lower monoids are incomparable.

Sections 3–5 prove the exhaustiveness assertion. Section 6 verifies the displayed inclusions and strictness.

3 Two-sided invariants and passage to an antichain

3.1 The preservation convention

For a structure Γ on a countable set D, $\mathrm{Emb}(\Gamma)$ denotes its injective maps $D\to D$ preserving and reflecting every basic relation. Reflection is equivalent to preserving the complement. These monoids are closed in D^D, with D discrete: failure of injectivity, preservation or reflection is witnessed on a finite tuple, so the set of maps with such a failure is open. A closed monoid of injections need not be an embedding monoid, so closedness alone will never replace the hypothesis $M=\mathrm{Emb}(\Gamma)$.

The existential preservation theorem for ω-categorical structures says that a relation is existentially definable without parameters precisely when it is preserved by all self-embeddings [9, Proposition 2 in the arXiv version]; see also [6]. Consequently, for two countable ω-categorical structures Γ_1,Γ_2 on the same domain D, equality $\mathrm{Emb}(\Gamma_1)=\mathrm{Emb}(\Gamma_2)$ is equivalent to equality of their collections of parameter-free existentially definable relations on D. In terms of signatures, this is equivalent to requiring, in both directions, that *each basic relation and its complement* in one structure be existentially definable without parameters in the other. Merely demanding existential definitions of the named positive relations would not justify the embedding-monoid correspondence. Theorem 2.3 gives five classes with the stated two-sided convention, or, equivalently, after complements have been named.

We say that a map h *violates* an n-ary relation T if some tuple $\bar{a}$ satisfies $T(\bar{a})$ but not $T(h\bar{a})$. This means failure of preservation; failure of reflection means instead that a false instance becomes true.

We will repeatedly use the following elementary observation.

Lemma 3.1 (Transfer of two-sided invariants)**.** *Let $M = \mathrm{Emb}(\Gamma)$ for a relational structure Γ on D. Let Θ be another relational structure on D, let $E \subseteq D$, and suppose:*

1. *there exists $g \in M \cap \mathrm{Emb}(\Theta)$ with $g[D] \subseteq E$;*
2. *every isomorphism between finite substructures of $\Theta|_E$ extends to a member of M.*

Then $\mathrm{Emb}(\Theta) \subseteq M$.

Proof. Fix $u \in \mathrm{Emb}(\Theta)$, a basic n-ary relation T of Γ, and $\bar{a} \in D^n$. The map $g(a_i) \mapsto g(u(a_i))$ is a well-defined isomorphism between finite substructures of $\Theta|_E$. Choose $h \in M$ extending it. Since g, h preserve and reflect T,

$$T(\bar{a}) \iff T(g\bar{a}) \iff T(hg\bar{a}) \iff T(gu\bar{a}) \iff T(u\bar{a}).$$

The middle equivalence uses the tuple identity $hg\bar{a} = gu\bar{a}$. This holds for every basic T and every tuple, so $u \in M$. □

Preservation and reflection give the same truth value for a basic relation on a tuple and on its image. We can therefore use each displayed equivalence in either direction, even though neither g nor h need have an inverse in M.

3.2 Published classification inputs

We state exactly the consequences of the earlier papers that we use.

Proposition 3.2 (Leaf classification [9, arXiv version, proof of Corollary 3 and Lemma 14])**.** *If Λ is a reduct of the homogeneous binary branching leaf structure $(L; C)$, then*

$$\mathrm{Emb}(\Lambda) \in \{\mathrm{Emb}(L; C),\ \mathrm{Emb}(L; Q),\ \mathrm{Inj}(L)\}.$$

Both $(L; C)$ and $(L; Q)$ are homogeneous. In each case the self-embedding monoid is the closure, in L^L, of the automorphism group, in accordance with the model-completeness criterion [11, Corollary 2.3].

For the monoid formulation, apply the cited proof to the expansion of Λ by the complements of its basic relations and by disequality. Its endomorphisms are then exactly the self-embeddings of Λ.

Proposition 3.3 (Flat-map reduction from [8])**.** *For a reduct Γ of Δ, one of the following holds:*

$$\mathrm{Emb}(\Gamma) = \mathrm{Emb}(S_2; \leq, C), \qquad \mathrm{Emb}(\Gamma) = \mathrm{Emb}(S_2; B), \qquad \mathrm{Emb}(\Gamma) = \mathrm{Inj}(S_2),$$

or $\mathrm{Emb}(\Gamma)$ *contains a flat self-embedding of $(S_2; R)$.*

Here is the precise dependency of Proposition 3.3. Write $M = \mathrm{Emb}(\Gamma)$. Expand Γ by the complements of its basic relations and by disequality; its endomorphisms are then exactly $\mathrm{Emb}(\Gamma)$. Proposition 4.17 of [8] gives the first two alternatives, or a flat or thin map, where a *thin* map has chain image. Lemmas 4.20 and 4.21 of that paper show that a thin self-embedding forces $\mathrm{Emb}(\Gamma) = \mathrm{Inj}(S_2)$. For a flat self-embedding, the distinct-triple

case analysis in the proof of Lemma 4.26 gives either the same equality alternative or a flat map preserving the injective relation (1). Since each distinct triple has exactly one rooted orientation, preserving R also reflects it. This is the flat-map consequence we use.

We spell out the outcomes of the cited flat-map case analysis that enter this reduction. Let $\prec$ be the homogeneous convex linear order from [8, Section 3.3]; its underlying order is isomorphic to $(\mathbb{Q};<)$. It extends $<$, puts a point incomparable with a comparable pair outside the interval spanned by the pair, and puts the outgroup of an antichain triple outside the interval spanned by the other two points. Here a map $g : S_2 \to S_2$ is *canonical* from $(S_2;\leq, C, \prec)$ to itself if, for every $n \geq 1$ and all $\bar{a}, \bar{b} \in S_2^n$, agreement of all quantifier-free formulas on $\bar{a}, \bar{b}$ in that structure implies agreement of all quantifier-free formulas on $g\bar{a}, g\bar{b}$. The image tuples must have the same type as each other; that type need not be their input type. Homogeneity makes this equivalent to the definition using complete first-order types. The proof of [8, Lemma 4.26] supplies such a flat map $g \in M$, normalised to preserve $\prec$. Appendix A verifies the Ramsey input needed for that canonical-map argument.

In two of the alternatives, the image outgroup of every triple $x \prec y \prec z$ is determined solely by its position in this order: it is always $g(z)$, or always $g(x)$. These are cases (A1)+(B3) and (A2)+(B1), respectively, in the cited proof. For arbitrary a, b, c, their complete descriptions in our injective convention are, respectively,

$$C(g(c); g(a), g(b)) \iff \mathrm{Dist}(c, a, b) \wedge a \prec c \wedge b \prec c,$$

and

$$C(g(c); g(a), g(b)) \iff \mathrm{Dist}(c, a, b) \wedge c \prec a \wedge c \prec b.$$

Fix a basic n-ary relation T of Γ and a quantifier-free $(<, C)$-formula φ defining it, which exists by quantifier elimination in Δ. Since $g \in \mathrm{Emb}(\Gamma)$, every $\bar{a} \in S_2^n$ satisfies

$$T(\bar{a}) \iff T(g\bar{a}) \iff \varphi(g\bar{a}).$$

The last expression denotes evaluation on the image tuple. To express its truth value by a formula on $\bar{a}$ in the original domain, replace every strict-order atom of φ by false, every equality atom by itself, and every C-atom by the appropriate displayed formula in $\prec$. Flatness and injectivity justify the first two replacements. Keeping the Boolean connectives produces a formula in $\prec$ and equality that defines T. Thus Γ is a reduct of $(S_2;\prec) \cong (\mathbb{Q};<)$.

Here the additional hypothesis that Γ is a reduct of the semilinear order is essential. Lemma 4.21 of [8] applies to a reduct of $(S_2;\leq)$ that is also isomorphic to a reduct of $(\mathbb{Q};<)$; its proof shows that every permutation is then an automorphism. Both hypotheses hold: C is definable from $\leq$, so Γ is a reduct of $(S_2;\leq)$, and the preceding formulas define it in $(S_2;\prec)$. Consequently each basic relation depends only on the equality pattern of its tuple. Every injection preserves and reflects these patterns, giving $M = \mathrm{Inj}(S_2)$.

In the remaining case, called (A1) together with (B4) in the cited proof, the action on distinct triples $x \prec y \prec z$ is as follows:

input type	image rooted orientation
$x < y < z$	$z' \mid x'y'$
$x \parallel y,\ x, y < z$	$z' \mid x'y'$
$x < y,\ z \parallel x,\ z \parallel y$	$z' \mid x'y'$
$x \parallel y,\ x \parallel z,\ y < z$	$x' \mid y'z'$
$C(z; x, y)$	$z' \mid x'y'$
$C(x; y, z)$	$x' \mid y'z'$

Here x', y', z' denote the images under g. The order conditions make the list exhaustive. The image outgroup in each row is exactly the outgroup prescribed by (1). This is the distinct-triple conclusion needed here.

Lemma 3.4. *In the last alternative of Proposition 3.3,* $\mathrm{Emb}(S_2; R) \subseteq \mathrm{Emb}(\Gamma)$.

Proof. Let $M = \mathrm{Emb}(\Gamma)$ and let $g \in M \cap \mathrm{Emb}(S_2; R)$ be flat. Choose a maximal antichain L containing its image. On L, $R = C$. Every finite C-isomorphism in L is a finite $(\leq, C)$-isomorphism in Δ, since all distinct points involved are incomparable. It extends to an automorphism of Δ, which belongs to M, because Γ is a reduct of Δ. Apply Lemma 3.1 with $\Theta = (S_2; R)$, $E = L$. □

3.3 Relations determined by leaf types

For $n \geq 1$, two tuples $\bar{a}, \bar{b} \in S_2^n$ have the same *R-type* if

$$a_i = a_j \iff b_i = b_j \quad \text{and} \quad R(a_i; a_j, a_k) \iff R(b_i; b_j, b_k)$$

for all $i, j, k \in \{1, \ldots, n\}$. Define equality of *$P$-types* by the same equality conditions and agreement of every corresponding P-instance on four coordinate indices. For tuples in L, equality of *C-types* and *Q-types* is defined by the equality pattern and agreement of every corresponding C- or Q-instance, respectively. A type thus records an ordered tuple. On a tuple of distinct entries, leaf labels are coordinate positions. If entries repeat, list the distinct entries in order of first occurrence as $c_1, \ldots, c_m$, and record the surjection $\pi : \{1, \ldots, n\} \to \{1, \ldots, m\}$ with $a_i = c_{\pi(i)}$; this retains every repeated position.

We recall why this relational information determines the representing tree. For distinct leaves y, z of a rooted binary leaf tree with leaf set A, put

$$D_{y,z} = \{y, z\} \cup \{x \in A \setminus \{y, z\} : \neg(x \mid yz)\}.$$

This is exactly the set of leaves below the least common ancestor of y, z: a leaf is outside that descendant set precisely when it is the outgroup to y, z. Every internal vertex is such a least common ancestor, by choosing one leaf in each child branch. Thus the sets $D_{y,z}$, together with the singleton leaf sets, recover all descendant sets. Different vertices have different descendant sets, and tree edges correspond to immediate inclusions among them. This reconstructs the labelled rooted tree, including the unique one- and two-leaf cases.

For an unrooted binary tree on at least four labelled leaves, choose a leaf p, delete it, and root the remaining tree at its former neighbour. On the other leaves, $x \mid yz$ holds exactly when the original quartet has split $px \mid yz$. Its quartet relation therefore determines the rooted tree, and reattaching p at the root recovers the original unrooted tree. On one, two or three labelled leaves there is a unique reduced unrooted tree: respectively a single vertex, an edge, or three edges meeting at one vertex. Here *reduced* means that no unlabelled vertex has degree two. Consequently R-types of injective tuples can be represented by labelled rooted trees, and P-types by labelled reduced unrooted trees.

Lemma 3.5. *Let Γ be a reduct of Δ, and suppose $M = \mathrm{Emb}(\Gamma)$ contains $\mathrm{Emb}(S_2; R)$. For every basic n-ary relation T of Γ, the following hold.*

1. *If $\bar{a}, \bar{b} \in S_2^n$ have the same R-type, then $T(\bar{a})$ holds if and only if $T(\bar{b})$ holds. Moreover, T is quantifier-free definable without parameters in $(S_2; R)$.*

2. *If M also contains $\mathrm{Emb}(S_2; P)$, then $T(\bar{a})$ holds if and only if $T(\bar{b})$ holds whenever $\bar{a}, \bar{b} \in S_2^n$ have the same P-type. In this case T is quantifier-free definable without parameters in $(S_2; P)$.*

Proof. Fix a basic n-ary relation T. For (1), choose a flat R-embedding f with image in a maximal antichain L. Let $\bar{a}, \bar{b} \in S_2^n$ have the same R-type. The tuples $f\bar{a}, f\bar{b}$ have the same

C-instances and equality pattern. All their distinct entries are incomparable, so the coordinate correspondence is a finite $(\leq, C)$-isomorphism. Homogeneity of Δ supplies $\alpha \in \mathrm{Aut}(\Delta) \subseteq M$ with $\alpha f\bar{a} = f\bar{b}$. Preservation and reflection under f, α give $T(\bar{a}) \leftrightarrow T(\bar{b})$.

For each realised n-tuple type, form the conjunction specifying the truth or falsity of every equality $x_i = x_j$ and every atom $R(x_i; x_j, x_k)$. There are finitely many such atoms and hence finitely many types. Disjoin these conjunctions for precisely those types whose realisations satisfy T. The equivalence just proved makes this selection well-defined, and the resulting quantifier-free formula defines T.

For (2), let $\bar{a}, \bar{b}$ have the same P-type and retain f above. Their image tuples have the same Q-instances and equality pattern. Homogeneity of $(L; Q)$, supplied by Proposition 3.2, gives $\beta \in \mathrm{Aut}(L; Q)$ with $\beta f\bar{a} = f\bar{b}$. The map $\beta f : S_2 \to L \subseteq S_2$ preserves and reflects P by the coordinatewise identity $P = f^{-1}(Q)$. Both f and βf belong to M, so

$$T(\bar{a}) \iff T(\beta f\bar{a}) \iff T(f\bar{b}) \iff T(\bar{b}).$$

Disjoining the complete equality and P-descriptions of the types where T holds gives the required quantifier-free definition.

These arguments include repeated coordinates. The common equality pattern makes the correspondence between distinct entries well-defined; an automorphism matching those entries matches every occurrence of them in the original tuples. □

Lemma 3.6 (Detecting a violation on an antichain)**.** *Let Γ be a reduct of Δ, and let $M = \mathrm{Emb}(\Gamma)$ contain $\mathrm{Emb}(S_2; R)$.*

1. *If a member of M violates R on an antichain triple, then $\mathrm{Emb}(S_2; P) \subseteq M$.*
2. *If $\mathrm{Emb}(S_2; P) \subseteq M$ and a member of M violates P on an antichain quadruple, then $M = \mathrm{Inj}(S_2)$.*

Proof. For (1), choose a maximal antichain L containing the violating input triple and a flat R-embedding f with image in L. By Lemma 3.5(1), every basic relation of Γ has a quantifier-free definition in R. Restricting that formula to L and replacing R by C shows that $\Gamma|_L$ is a quantifier-free reduct of $(L; C)$. If $h \in M$ is the violating map, $(fh)|_L$ is an embedding of $\Gamma|_L$ that violates C. Proposition 3.2 therefore gives $\mathrm{Emb}(L; Q) \subseteq \mathrm{Emb}(\Gamma|_L)$. Homogeneity of $(L; Q)$ sends any tuple to any other tuple of the same Q-type by an automorphism. Such automorphisms preserve every basic relation of $\Gamma|_L$, so each of those relations has the same truth value on the two tuples. If $\bar{a}, \bar{b}$ have the same P-type, $f\bar{a}, f\bar{b}$ have the same Q-type; two-sided preservation under f shows that each relation of Γ has the same value on $\bar{a}, \bar{b}$. Every P-embedding consequently belongs to M.

For (2), let $h \in M$ witness the assumed violation of P, and choose L, f using its violating quadruple as in part (1). Lemma 3.5(2) gives quantifier-free definitions in P. On L, these become quantifier-free definitions in Q, so every Q-embedding preserves and reflects the basic relations of $\Gamma|_L$. Its embedding monoid therefore contains $\mathrm{Emb}(L; Q)$, as well as the embedding $(fh)|_L$ violating Q. Proposition 3.2 makes that monoid $\mathrm{Inj}(L)$. Its relations depend only on equality patterns. Pulling them back along f shows the same for all relations of Γ, so every injection of S_2 belongs to M. □

4 Rooted type transitions

Lemma 3.6 handles violations witnessed on antichains. We now consider a map whose violation of R requires comparable points. The aim is to show that the basic relations of the reduct then depend only on equality. We shall associate a finite graph to each tuple

length: its vertices are rooted tree types, and an edge records a change made by the given map. For each basic relation, tuples realising the two endpoint types have equal truth values: Lemma 3.5 makes the choice of realisation immaterial, and the witnessing map preserves and reflects the relation. The proof therefore reduces to connecting all vertices by a few local tree moves. The one-map formulation below will also be used for P in Section 5.

Theorem 4.1 (Rooted reconstruction)**.** *Let $M = \mathrm{Emb}(\Gamma) \supseteq \mathrm{Emb}(S_2; R)$, where Γ is a reduct of Δ. Suppose that M contains a flat map k which preserves C on every antichain and violates R on some triple. Then $M = \mathrm{Inj}(S_2)$.*

4.1 Reduction to a branching witness

There are three kinds of true R-triples outside antichains:

$$\begin{aligned} R_1(x;y,z): &\quad x \parallel y,\ x \parallel z,\quad y < z \text{ or } z < y,\\ R_2(x;y,z): &\quad y, z < x,\quad y \parallel z,\\ R_3(x;y,z): &\quad y < z < x \text{ or } z < y < x. \end{aligned}$$

We say that a map preserves R on R_i if the image of every triple in R_i satisfies R; the image need not still have order pattern R_i.

Lemma 4.2. *A flat map k preserving C on every antichain preserves R on R_1. If it also preserves R on R_2, it preserves R everywhere. Consequently, under the hypotheses of Theorem 4.1, there exist $a, b, c \in S_2$ such that*

$$b, c < a, \qquad b \parallel c, \qquad C(k(b); k(a), k(c)).$$

Proof. Consider $x \parallel y$ and $z < y$. Niceness gives $d > y$ with $d \parallel x$, and binary branching below d gives $w < d$ with $w \parallel y$. Then $w \parallel z$, and the antichain triples (x, y, w) and (x, z, w) satisfy C. Their images satisfy the same two rooted orientations. Choose a finite rooted leaf tree representing these image points, using Lemma 2.2, and label its leaves by the points they represent. In this tree the relation

$$u \equiv_t v \quad \Longleftrightarrow \quad u = v \ \text{ or } \ t \mid uv \qquad (u, v \neq t)$$

is an equivalence relation: each class is the set of leaves in one of the subtrees branching off the path from the root to t. Both $k(y)$ and $k(z)$ are equivalent to $k(w)$ for $t = k(x)$; hence $C(k(x); k(y), k(z))$. The case $y < z$ is symmetric.

It remains to consider $z < y < x$, assuming preservation on R_2. Choose $w < x$ with $w \parallel y$. The already proved R_1 case and the R_2 hypothesis give

$$C(k(w); k(y), k(z)), \qquad C(k(x); k(y), k(w)).$$

Represent these four image points by a rooted leaf tree. In the following comparison, *below* refers to its ancestor order. The latter relation places $k(y), k(w)$ together below their separation from $k(x)$. The former places $k(z)$ on the $k(y)$-side of the separation of $k(y)$ from $k(w)$. Thus $k(z)$ lies on the same side as $k(y), k(w)$ relative to $k(x)$, and $C(k(x); k(y), k(z))$ follows. This verifies R_3.

The four forms C, R_1, R_2, R_3 exhaust R. A violation must therefore be on R_2. On its image exactly one rooted orientation holds; exchanging b, c if necessary gives the displayed orientation. □

4.2 The finite type graphs and their local moves

Fix k and a, b, c as in Lemma 4.2. For each $n \geq 1$, let $\mathcal{T}_n$ be the finite set of label-preserving isomorphism classes of rooted binary leaf trees with leaf labels $1, \ldots, n$. Internal vertex names are immaterial. Define a directed graph G_n on $\mathcal{T}_n$ by placing an edge $s \to t$ whenever an injective tuple $\bar{x} \in S_2^n$ has R-type s and $k\bar{x}$ has R-type t. We use connectedness in the underlying undirected graph. Reversing a graph edge does not mean inverting k. We also use G_n on any specified n-element label set via a bijection with $\{1, \ldots, n\}$; changing that bijection merely relabels the graph.

Fix a basic n-ary relation T of Γ. Define $\varepsilon_T : \mathcal{T}_n \to \{0, 1\}$ by setting $\varepsilon_T(s) = 1$ if T holds on an injective tuple of R-type s, and $\varepsilon_T(s) = 0$ otherwise. Every s is realised by Lemma 2.2, and Lemma 3.5 makes the value independent of the realising tuple. A tuple witnessing $s \to t$ shows $\varepsilon_T(s) = \varepsilon_T(t)$, because k preserves and reflects T. Thus this function has one value on each connected component. This graph argument concerns injective tuples; a coordinate map will handle repeated entries at the end of the section. Relabelling the coordinates is an automorphism of G_n: the same witnessing tuple with its coordinates permuted witnesses the relabelled edge.

If s, t are trees on disjoint sets of labels, write $s \star t$ for the tree formed by attaching their roots to a new common root. We write $\{a\}$ for the one-leaf rooted tree labelled a.

Lemma 4.3 (Adding a branch to a transition)**.** *If s, s' are connected in G_m, and t is a rooted binary leaf tree on $r \geq 1$ further labels, then $s \star t$ and $s' \star t$ are connected in G_{m+r}.*

Proof. We shall show that every directed edge $s \to s'$ in G_m gives a directed edge $s \star t \to s' \star t$ in G_{m+r}. In other words, a change of tree type made by k can still occur after an additional branch of type t has been attached. The additional points need not be fixed by k; we must ensure that their tree type is unchanged and that their images remain a separate branch.

Let $\bar{x} = (x_1, \ldots, x_m)$ witness $s \to s'$, so that $\bar{x}$ has R-type s and $k\bar{x}$ has R-type s'. Put $A = \{x_1, \ldots, x_m\}$ and choose a strict common upper bound d of A. Take a normalised diagram of $A \cup \{d\}$, and attach its root and the root of t below a new common root. Lemma 2.1, with $A \cup \{d\}$ fixed, realises this larger diagram. Its new marked leaves form an antichain F, labelled as t, every point of which is incomparable with d. List them as a tuple $\bar{y}$ according to the chosen identification of their labels with coordinate positions. The combined tuple $(\bar{x}, \bar{y})$ has R-type $s \star t$.

We check its image type. On the A-coordinates the image type is s', by the choice of $\bar{x}$. On F it is t, because k preserves C on antichains. For two distinct points $x, x' \in A$ and $y \in F$, the relation $R(y; x, x')$ is either C, when x, x' are incomparable, or of form R_1, when they are comparable. It is preserved by the hypothesis on k and Lemma 4.2. For $x \in A$ and distinct $y, y' \in F$, the relation $C(x; y, y')$ holds and is also preserved. These are all triples meeting both parts, and their orientations say exactly that the two parts occupy separate branches. Thus the image has R-type $s' \star t$, proving the required edge assertion. If one part is a singleton, only the mixed triples with two entries in the other part need to be checked; if both are singletons, there is only one two-leaf tree.

Finally, take an undirected path from s to s' in G_m. Applying the edge assertion to each of its edges, and traversing the resulting edges in the corresponding directions, gives a path from $s \star t$ to $s' \star t$. □

A *cherry* is a pair of leaves with the same parent. A *pendant subtree* at a vertex consists of that vertex and all its descendants. A pendant subtree with three leaves consists of a cherry and its sibling leaf.

Lemma 4.4 (Moving a leaf to the root)**.** *Let T be a labelled rooted binary leaf tree containing a pendant three-leaf subtree $b \mid ac$, with a, c its cherry. Let $F = T \setminus \{a\}$, suppressing the*

resulting vertex with one child. If T has n leaves, then T and $\{a\} \star F$ are connected in G_n on that same label set.

Proof. The letters a, b, c in this statement are leaf labels. For the realisation below, assign them respectively to the three points in the fixed violating witness: a is the upper point, and the image of b is the outgroup. This assignment identifies labels with points only for this construction. In F, the leaves b, c form a cherry. We realise F as an antichain below the point a in the fixed violating triple from Lemma 4.2, with its leaves b, c at those fixed points. Here is why this simultaneous choice is possible. Take the full finite tree F, add a marked point a strictly above its root, and retain b, c as marked leaves. Its restriction to these three marked points has exactly the type $b, c < a,\ b \parallel c$. Lemma 2.1 therefore realises the entire diagram over the already fixed a, b, c.

The R-type of the resulting tuple is $\{a\} \star F$. Its F-coordinates form an antichain, so their images still have tree F. Write $a' = k(a)$, $b' = k(b)$, and $c' = k(c)$, and let T' be the tree of the entire image tuple. Deleting a' from T' and suppressing the resulting vertex with one child recovers F, with each label placed at its image point. Hence T' is obtained from F by subdividing an edge and attaching the new leaf a' there, or by putting a' beside F under a new root.

The witness $C(b'; a', c')$ determines which attachment is possible. It requires the least common ancestor of a', c' to lie strictly below the least common ancestor of b', c'. In the tree F recovered by deleting a' from T' and suppressing the resulting vertex with one child, the leaves b', c' form a cherry: from their common parent down towards c' there is just the edge ending at c'. Consequently the new leaf must be attached by subdividing that edge. Attaching it anywhere else, including above the root, would fail the required rooted orientation. The resulting three-leaf subtree has outgroup b' and cherry a', c', while the rest of F is unchanged. Thus T' is exactly the labelled tree T. This gives an edge from $\{a\} \star F$ to T, and hence a connection in either direction. □

A *comb* is a rooted binary leaf tree in which every non-leaf vertex has a leaf child. For at least three leaves its *top leaf* is the singleton branch at the root. A *balanced four-leaf tree* has two cherries as its root branches.

Lemma 4.5 (The four-leaf transition)**.** *On the label set $\{a, b, c, d\}$, the graph G_4 has a directed edge from the comb $\{a\} \star (\{c\} \star (\{b\} \star \{d\}))$ to the balanced tree $(\{a\} \star \{c\}) \star (\{b\} \star \{d\})$. In particular, these two vertices are connected.*

Proof. For the fixed $b, c < a$ witness, choose $d < b$. The input R-tree is the comb

$$\{a\} \star (\{c\} \star (\{b\} \star \{d\})) :$$

its top leaf is a, the next leaf is c, and its bottom cherry is $\{b, d\}$. Write $a' = k(a)$, and similarly for the other points. We have

$$C(b'; a', c') \quad \text{and} \quad C(c'; b', d').$$

The first is the violating witness; the second follows from preservation of $R_1(c; b, d)$. Let v be the least common ancestor of a', c' in the image tree, and let w be the least common ancestor of b', d'. The first relation says that b' is not a descendant of v; the second says that c' is not a descendant of w. Descendant sets of vertices in a rooted tree are either nested or disjoint. The set below v cannot contain the set below w, because it excludes b', and the reverse containment is impossible because the set below w excludes c'. Thus these two sets are disjoint. They contain respectively a', c' and b', d', accounting for all four leaves. As the tree is binary, they are its two root branches. The image is therefore the balanced tree with root split $\{a', c'\} \mid \{b', d'\}$. Figure 5 shows this transition. □

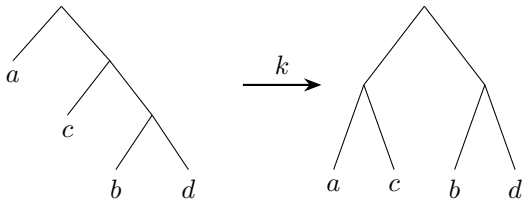


Figure 5: The R-trees of the source and image in Lemma 4.5. The source is realised by $d < b < a$, $c < a$, $c \parallel b$. The drawing shows leaf types, not the original order of these four points.

4.3 Connectivity and reconstruction

Proposition 4.6. *For every $n \geq 1$, the graph G_n is connected.*

Proof. We induct on the number n of labelled leaves. The induction assertion is connectivity for every choice of n distinct labels; relabelling makes this equivalent to connectivity of G_n. For $n = 1, 2$ there is one vertex. For $n = 3$, the fixed violating triple gives an edge between two different outgroups. Permuting the three labels gives edges connecting all three possible outgroups.

We first record how the induction hypothesis controls labels on combs. Let $n \geq 4$, suppose connectivity is known in all smaller arities, and let H be a component containing a comb with top leaf p. By Lemma 4.3 and connectivity of G_{n-1}, we may replace its lower tree by any tree on the other $n - 1$ labels. In particular, every relabelling fixing p maps this comb to a vertex of H. A graph automorphism sends a connected component to a connected component; since the image component intersects H, it equals H. Thus the whole point stabiliser $\mathrm{Sym}(\{1, \ldots, n\})_p$ preserves H.

The comb has a pendant three-leaf subtree at its bottom. Choose one cherry leaf q of that subtree. Lemma 4.4 moves q to the root. Connectivity in arity $n - 1$ and Lemma 4.3 then turn the remaining tree into a comb. Hence H also contains a comb with top leaf $q \neq p$, and the same argument gives its entire point stabiliser. These two stabilisers generate $\mathrm{Sym}(n)$. Indeed, every transposition except possibly $(p\,q)$ fixes p or q; if $r \notin \{p, q\}$, then $(p\,q) = (p\,r)(q\,r)(p\,r)$. All rooted binary combs with n leaves have the same unlabelled shape: successively removing the top leaf leaves a smaller comb, ending in a two-leaf cherry. Every labelled comb is therefore a relabelling of the one already in H. The invariance under $\mathrm{Sym}(n)$ just proved puts every labelled comb in H.

For $n = 4$, every rooted binary tree is either a comb or balanced. The preceding argument puts all labelled combs in one component. Lemma 4.5 adds a balanced tree to it, and invariance under $\mathrm{Sym}(4)$ adds every labelling of the balanced tree. This proves the base case required beyond arity three.

Now let $n \geq 5$ and take any tree T. Each of its two root branches has fewer than n leaves. Using the induction hypothesis and Lemma 4.3, replace each branch by a comb on the same label set. At least one branch has at least three leaves. Its bottom three leaves form a pendant three-leaf subtree, so Lemma 4.4 moves one of its cherry leaves to the root of the whole tree. The remaining $n - 1$-leaf tree can then be replaced by a comb. Thus every component contains a comb. As already proved, each component containing a comb contains all labelled combs. Any two such components intersect and are equal, proving connectivity. □

Proof of Theorem 4.1. Let T be a basic n-ary relation of Γ, let $1 \leq m \leq n$, and let $\pi : \{1, \ldots, n\} \to \{1, \ldots, m\}$ be a surjection. Define

$$T_\pi(x_1, \ldots, x_m) \quad \Longleftrightarrow \quad T(x_{\pi(1)}, \ldots, x_{\pi(n)}).$$

On injective m-tuples, π specifies a fixed equality pattern for the original n coordinates. Two m-tuples with the same R-type give two n-tuples with the same R-type after their coordinates are repeated according to the same map π. Lemma 3.5, applied to T, thus gives equal truth values for T_π on the two m-tuples. Also, k preserves and reflects T_π, since its coordinatewise application commutes with that repetition of coordinates. Hence T_π defines a function $\varepsilon_{T_\pi} : \mathcal{T}_m \to \{0,1\}$, by the same rule as above, with equal values at the ends of every edge of G_m. Proposition 4.6 makes this function constant. Consequently T_π has the same truth value on every injective m-tuple. Varying m, π shows that T depends only on its equality pattern. Every injection preserves and reflects such relations, so $M = \mathrm{Inj}(S_2)$. □

Corollary 4.7. *Let Γ be a reduct of Δ, and let $M = \mathrm{Emb}(\Gamma) \supseteq \mathrm{Emb}(S_2; R)$. If every member of M preserves R on every antichain triple, then $M = \mathrm{Emb}(S_2; R)$. Consequently, if this inclusion is strict, then $\mathrm{Emb}(S_2; P) \subseteq M$.*

Proof. If $h \in M \setminus \mathrm{Emb}(S_2; R)$, then h violates R: for distinct triples the three rooted orientations are exhaustive, so preserving R would also reflect it. Compose h on the left with a flat R-embedding f. The map $k = fh$ still violates R and, under the hypothesis, preserves C on antichains. Theorem 4.1 gives $M = \mathrm{Inj}(S_2)$, contrary to preservation of R on antichains. This proves the first assertion. If the inclusion is strict, it therefore has an antichain violation, to which Lemma 3.6(1) applies. □

5 Unrooted quartets and an auxiliary root leaf

We next prove that no embedding monoid of a reduct of Δ lies strictly between $\mathrm{Emb}(S_2; P)$ and $\mathrm{Inj}(S_2)$. The difficulty is again a violation away from antichains. Fixing an auxiliary leaf converts such a violation into the rooted situation already settled.

Lemma 5.1 (Rooting at a leaf)**.** *Let $(L; Q)$ be the homogeneous binary branching quartet structure and $p \in L$. On $L \setminus \{p\}$ define, for distinct x, y, z,*

$$C_p(x; y, z) \quad \Longleftrightarrow \quad Q(p, x; y, z).$$

Set C_p to false on repeated tuples. Then $(L \setminus \{p\}; C_p)$ is isomorphic to $(L; C)$. The relation defined on distinct quadruples by formula (2) with C_p replacing R, and set to false on repeated quadruples, is the original Q restricted to $L \setminus \{p\}$.

Proof. In a finite unrooted binary tree containing p, remove the leaf p and root the remaining tree at its former neighbour, suppressing unnecessary vertices. The rooted orientation $x \mid yz$ in this tree is exactly the quartet split $px \mid yz$ before removing p. Forgetting that root recovers all quartet splits on the remaining leaves. This proves the last assertion and shows that every finite induced C_p-structure is the relational structure on the leaves of a rooted binary leaf tree.

Conversely, attach a new leaf p above the root of any finite rooted binary leaf tree and forget the root. The resulting finite quartet structure embeds into $(L; Q)$, by its universality for finite binary leaf structures. Choose such an embedding. Homogeneity of $(L; Q)$ gives an automorphism taking the image of its distinguished leaf to the specified p. Composing with that automorphism gives the desired rooted leaf structure in $(L \setminus \{p\}; C_p)$.

Finally, a finite C_p-isomorphism extends, by fixing p, to a finite Q-isomorphism. Quartets containing p are determined by the defining formula, and quartets not containing p are determined by the recovered quartet relation just proved. Homogeneity of $(L; Q)$ extends this isomorphism to an automorphism fixing p. Hence $(L \setminus \{p\}; C_p)$ is homogeneous and has exactly the age of $(L; C)$. The uniqueness part of Fraïssé's theorem—two countable homogeneous

structures with the same age are isomorphic—therefore gives the claimed isomorphism. This is the finite-tree rerooting construction also used in [9, arXiv version, proof of Lemma 14]. □

Lemma 5.2 (Interpolation on an antichain)**.** *Let L be a maximal antichain in S_2. Every isomorphism between finite Q-substructures of L extends to a member of* $\mathrm{Emb}(S_2; P)$.

Proof. Let $u : A \to B$ be the given map, with $A, B \subseteq L$. By Lemma 2.2, choose a flat R-embedding $f : S_2 \to L$ fixing A. Homogeneity of $(L; Q)$ extends u to $\beta \in \mathrm{Aut}(L; Q)$. Then βf extends u and preserves and reflects $P = f^{-1}(Q)$. □

Theorem 5.3 (Quartet reconstruction)**.** *If $M = \mathrm{Emb}(\Gamma)$ is an embedding monoid of a reduct of Δ and*

$$\mathrm{Emb}(S_2; P) \subsetneq M,$$

then $M = \mathrm{Inj}(S_2)$.

Proof. If some member of M violates P on an antichain, Lemma 3.6(2) finishes the proof. Suppose, therefore, that every member of M preserves P on all antichains. We shall construct a map to which Theorem 4.1 applies. The construction must retain a quartet violation while preserving rooted triples on antichains; an auxiliary leaf will specify the root in the image.

Choose $h \in M \setminus \mathrm{Emb}(S_2; P)$. The three quartet splits are exhaustive on distinct quadruples, so h violates P, rather than merely its complement. Composing on the left with a flat R-embedding makes h flat and retains this violation. Fix a violating ordered quadruple $\bar{a} = (a_1, a_2, a_3, a_4)$, and let $A = \{a_1, a_2, a_3, a_4\}$ be its set of entries. Choose a strict common upper bound d of A and a point $r \parallel d$.

The strict lower cone $U = \{x : x < d\}$, with its induced order and C-relation, is isomorphic to Δ. To check this, the density, binary branching and niceness properties listed in Section 2 restrict to U; a common upper bound within U follows from the no-joins property, and an upper point within U follows from density below d. There are no endpoints in U. A niceness witness, or a binary-branching witness for an antichain triple in U, is automatically below d: it is comparable with d because they have a common lower point, and it cannot be above d while remaining incomparable with another point below d. The uniqueness assertion of [8, Theorem 3.5] therefore applies. Choose an isomorphism $e : \Delta \to \Delta|_U$, regarded also as a self-embedding of Δ. Then $e \in M$, and $A \subseteq e[S_2]$.

Choose a maximal antichain L containing $h[S_2]$, and write $p = h(r)$. Lemma 5.1 gives an isomorphism

$$\iota : (L \setminus \{p\}; C_p) \longrightarrow (L; C).$$

By the last assertion of Lemma 5.1 and the quartet-defining formula, ι also preserves and reflects Q between its induced structures on $L \setminus \{p\}$ and L. Define $k = \iota h e : S_2 \to L$, and regard it as a self-map of S_2 through $L \subseteq S_2$. The composition is defined because $r \notin U$ and injectivity of h gives $he[S_2] \subseteq L \setminus \{p\}$.

We verify that $k \in M$. On any finite set $F \subseteq S_2$, the map $\iota|_{he[F]}$ is a finite Q-isomorphism inside L. By Lemma 5.2, it agrees there with some $v_F \in \mathrm{Emb}(S_2; P) \subseteq M$. Thus $v_F h e \in M$ agrees with k on F. Since an embedding monoid is closed in the topology of pointwise convergence, $k \in M$.

For an antichain triple x, y, z satisfying $C(x; y, z)$, the four points $r, e(x), e(y), e(z)$ form an antichain. Since e preserves C, we have $C(e(x); e(y), e(z))$. Also d is a common strict upper bound of $e(y), e(z)$ and is incomparable with r, so $C(r; e(y), e(z))$. The definition (2) therefore gives

$$P(r, e(x); e(y), e(z)).$$

By our assumption on antichains, the h-image satisfies P. All image entries lie in L, where $P|_{L^4} = Q$, so it satisfies $Q(p, he(x); he(y), he(z))$, or equivalently $C_p(he(x); he(y), he(z))$. Applying ι gives $C(k(x); k(y), k(z))$. Hence k preserves C on every antichain.

On the other hand, each a_i lies in the image of e, so the tuple $e^{-1}\bar{a} = (e^{-1}(a_1), \ldots, e^{-1}(a_4))$ is defined. The map k violates P on this tuple: e preserves P, h violates it on $\bar{a}$, and ι preserves and reflects the quartet relation. Since P is given by (2), a map preserving and reflecting R would preserve P. Thus k violates R. It is flat by construction. Theorem 4.1 now yields $M = \mathrm{Inj}(S_2)$, as required. □

6 Completion of the classification

Proof of the exhaustiveness assertion in Theorem 2.3. Let $M = \mathrm{Emb}(\Gamma)$. Proposition 3.3 either gives one of the three monoids listed there, or supplies a flat R-embedding in M. In the latter case Lemma 3.4 gives $\mathrm{Emb}(S_2; R) \subseteq M$. If this inclusion is strict, Corollary 4.7 gives $\mathrm{Emb}(S_2; P) \subseteq M$. If that inclusion is also strict, Theorem 5.3 gives $M = \mathrm{Inj}(S_2)$. All five displayed monoids are embedding monoids of reducts because their defining relations were explicitly defined in Δ. □

We finish by verifying the inclusion diagram.

Lemma 6.1. *Every self-embedding of $(S_2; B)$ preserves and reflects P.*

Proof. For pairwise distinct a, b, c, d, the following formula expresses P:

$$\exists x\, [B(a, x, c) \wedge B(a, x, d) \wedge B(b, x, c) \wedge B(b, x, d)]. \tag{3}$$

Suppose first that the leafified quartet has split $ab \mid cd$. In the unrooted tree, let v be the branching vertex at which the branches towards a, b meet the rest of the tree, and let w be the corresponding vertex for c, d. The segment from v to w belongs to each of the four paths joining a to c, a to d, b to c, and b to d. An interior point x of this segment would therefore satisfy all four displayed betweenness conditions.

We must realise such a point in S_2, rather than only in the leaf tree. Before suppressing vertices, the segment contains an edge of the original marked diagram: every edge added during leafification ends at a newly attached leaf, so it cannot separate two pairs of leaves. Insert a new marked point in the interior of an original edge on the segment. Lemma 2.1, applied over the fixed points a, b, c, d, realises the enlarged diagram in S_2. The path interpretation of B gives the four required instances with this point as x.

Conversely, suppose that a point x satisfies the four betweenness conditions, and take a diagram of a, b, c, d, x. If the quartet split were $ac \mid bd$, the paths from a to c and from b to d would be disjoint. If it were $ad \mid bc$, the paths from a to d and from b to c would be disjoint. Neither case permits x to lie on all four paths. Under leafification, each path between original marked points is contained in the corresponding path between the leaves receiving their labels: the operation only adds terminal edges to that path. Disjointness of the leaf paths therefore implies disjointness of the original paths. This proves the equivalence with (3).

A B-embedding preserves (3), since it maps a witness to a witness. It preserves the other two quartet splits by the same argument with permuted variables. Exactly one split holds on distinct quadruples; hence it also reflects P. Injectivity handles the false repeated-coordinate cases. □

Proof of the inclusion and strictness assertions in Theorem 2.3. The defining formulas give

$$\mathrm{Emb}(\Delta) \subseteq \mathrm{Emb}(S_2; R) \subseteq \mathrm{Emb}(S_2; P).$$

We also have $\mathrm{Emb}(\Delta) \subseteq \mathrm{Emb}(S_2; B)$, and Lemma 6.1 gives $\mathrm{Emb}(S_2; B) \subseteq \mathrm{Emb}(S_2; P)$.

A flat R-embedding fails to preserve $<$; it also sends a three-point chain, on which a betweenness instance holds, to an antichain, on which no betweenness instance holds. Thus $\mathrm{Emb}(S_2; R) \not\subseteq \mathrm{Emb}(S_2; B)$. In the other direction, [8, Definition 4.2 and Proposition 4.3] supplies a rerooting in $\mathrm{Emb}(S_2; B)$ which reverses a nonempty upward-closed chain. Such a chain contains $a < b < c$. The true instance $R(c; a, b)$ becomes false after reversal, because the lowest image point is then the image of c, whereas the outgroup of a chain is its highest point. Consequently $\mathrm{Emb}(S_2; B) \not\subseteq \mathrm{Emb}(S_2; R)$. These two witnesses prove all four strict inclusions below $\mathrm{Emb}(S_2; P)$ and the incomparability of the two lower middle monoids.

Finally choose an antichain quartet with split $ab \mid cd$. A permutation exchanging b, c and fixing a, d violates this split, so $\mathrm{Emb}(S_2; P) \subsetneq \mathrm{Inj}(S_2)$. Exhaustiveness already proved excludes any additional monoids or inclusions. □

7 Failure of existential finiteness

The finite classification in Theorem 2.3 does not extend to all homogeneous structures in finite relational languages. We give an equivalence-relation example with exactly five reducts up to first-order interdefinability but infinitely many distinct self-embedding monoids of reducts. A homogeneous Ramsey expansion of the same example shows that imposing the Ramsey property does not restore existential finiteness. This answers negatively the proposed existential strengthening discussed in [37, Conjecture 4.5] and [38].

7.1 An equivalence relation with singleton and infinite classes

Let

$$D = S \sqcup \bigsqcup_{i \in \mathbb{N}} C_i, \qquad S = \{s_0, s_1, \ldots\},$$

where S and every C_i are countably infinite and these sets are pairwise disjoint. Let E be the equivalence relation whose classes are the C_i and the singletons $\{s_j\}$, and put $U = D \setminus S$. We regard U as a unary relation and write

$$\Omega = (D; E, U).$$

This is a homogeneous structure in a finite relational language. Indeed, a finite partial isomorphism maps points of S to points of S, and matches the finitely many infinite classes that it meets. Extend its map on S to a permutation of S, its matching of infinite classes to a permutation of their countable index set, and its partial maps within matched classes to bijections. The resulting permutation of D is an automorphism extending the given partial isomorphism. Also

$$U(x) \iff \exists y\, (y \neq x \wedge E(x, y)), \tag{4}$$

so Ω and $(D; E)$ are first-order interdefinable. The latter structure is not homogeneous: for $c \in C_0$, the one-point map $s_0 \mapsto c$ is a partial isomorphism, but no automorphism can extend it, since automorphisms preserve equivalence-class cardinalities.

For each integer $n \geq 2$, define

$$R_n(x_1, \ldots, x_n) \iff \bigwedge_{i<j} x_i \neq x_j \ \wedge \ \bigwedge_{i=2}^{n} E(x_1, x_i), \qquad S_n = D^n \setminus R_n,$$

and let $\Gamma_n = (D; S_n)$. Here S_n is a relation, whereas S is the set of singleton-class points. Each Γ_n is a quantifier-free reduct of Ω with one basic relation. Since embeddings preserve and reflect their basic relations, $\mathrm{Emb}(\Gamma_n) = \mathrm{Emb}(D; R_n)$.

Theorem 7.1. *The monoids* $\mathrm{Emb}(\Gamma_n)$*,* $n \geq 2$*, form the strictly increasing chain*

$$\mathrm{Emb}(\Gamma_2) \subsetneq \mathrm{Emb}(\Gamma_3) \subsetneq \mathrm{Emb}(\Gamma_4) \subsetneq \cdots .$$

For $2 \leq n < m$*, the named relation* S_n *has no existential definition in* Γ_m*. Nevertheless every* Γ_n *is first-order interdefinable with* $(D; E)$*.*

Proof. We first prove inclusion and then construct an embedding that witnesses strictness. For every $n \geq 2$,

$$R_{n+1}(x_1, \ldots, x_{n+1}) \iff \bigwedge_{j=1}^{n+1} R_n(x_1, \ldots, \widehat{x_j}, \ldots, x_{n+1}).$$

The forward implication follows from the definition. Conversely, every pair of coordinates occurs together in one of the displayed n-subtuples, so the right-hand side forces all coordinates to be distinct and E-equivalent. Thus every embedding preserving and reflecting R_n also preserves and reflects R_{n+1}.

Fix $n \geq 2$. Choose bijections $b_i : C_i \to C_{i+1}$ for all $i \in \mathbb{N}$, and distinct points $c_0, \ldots, c_{n-1} \in C_0$. Define an injection $f_n : D \to D$ by

$$f_n(x) = b_i(x) \quad (x \in C_i), \qquad f_n(s_j) = \begin{cases} c_j, & j < n, \\ s_{j-n}, & j \geq n. \end{cases}$$

The images of the specified pieces are disjoint. Let $m > n$. Every R_m-tuple lies in one infinite source class and is mapped into one infinite target class, so f_n preserves R_m. To check reflection, an R_m-tuple in the image cannot lie in C_0, which contains only n image points. If it lies in C_{i+1}, all of its preimages lie in C_i. Its coordinates have distinct preimages, so these preimages form an R_m-tuple. Consequently $f_n \in \mathrm{Emb}(\Gamma_m)$.

On the other hand, $(s_0, \ldots, s_{n-1})$ belongs to S_n, while its image does not. Hence $f_n \notin \mathrm{Emb}(\Gamma_n)$, proving every strict inclusion. The same witness proves the assertion about existential definitions: an embedding preserves an existential formula by transporting its witnesses, since it preserves all quantifier-free formulas, including negated atoms. If S_n had an existential definition in Γ_m, then f_n would preserve S_n.

Finally,

$$E(x, y) \iff x = y \ \vee \ \exists z_3 \cdots z_n \, \neg S_n(x, y, z_3, \ldots, z_n).$$

For $n = 2$ the quantifier block is empty. For $n > 2$, two distinct E-equivalent points lie in an infinite class and can be extended to n distinct points there. This proves the displayed definition in both directions. Conversely, S_n is quantifier-free definable from E, so the structures are first-order interdefinable. □

In particular, all the reducts Γ_n belong to one first-order interdefinability class. Their embedding monoids have the same group of invertible elements, namely $\mathrm{Aut}(D; E)$. The complement in the definition of Γ_n also makes the existential separation visible on a named relation. This is consistent with the two-sided convention of Section 3: although E is existentially definable in each Γ_n, its complement is not existentially definable there for $n \geq 3$. Indeed, $f_2 \in \mathrm{Emb}(\Gamma_n)$ sends the non-E-related pair (s_0, s_1) into C_0. For $n = 2$, both E and its complement are quantifier-free definable. The theorem proves failure of existential finiteness even without appealing to a preservation theorem.

7.2 The five first-order reducts of the unordered structure

Define two further equivalence relations on D by

$$F(x,y) \iff E(x,y) \vee (x \in S \wedge y \in S), \qquad J(x,y) \iff (U(x) \leftrightarrow U(y)).$$

The F-classes are $S, C_0, C_1, \ldots$, whereas the J-classes are S, U.

Theorem 7.2. *Up to first-order interdefinability, the reducts of Ω are exactly*

$$(D;E), \qquad (D;F), \qquad (D;U), \qquad (D;J), \qquad (D;=).$$

Writing G_E, G_F, G_U, G_J for their respective automorphism groups, their strict inclusions are precisely the two chains

$$G_E < G_U < G_J < \mathrm{Sym}(D), \qquad G_E < G_F < \mathrm{Sym}(D),$$

with G_F incomparable with both G_U and G_J.

Finiteness also follows from [16, Theorems 1.2 and 5.11]. The relevant construction is the operation of taking countably many copies [16, Definition 2.29]: from a structure on A, it forms a structure on $A \times \mathbb{N}$, puts the original relations separately on each $A \times \{i\}$, and adds the equivalence relation of belonging to the same copy. Applied to a one-element structure, it gives a countable pure set up to first-order interdefinability. Applied again, it gives infinitely many infinite equivalence classes. Taking the disjoint union of this structure with a countable pure set, with a unary predicate naming each summand [16, Definition 2.28], gives Ω up to first-order interdefinability. These are construction operations covered by the cited theorem. The following proof determines the five reducts explicitly.

Proof. Put $G = G_E = \mathrm{Aut}(\Omega)$. Its elements act arbitrarily on S, permute the infinite classes arbitrarily, and act by arbitrary bijections between those classes. In particular, G contains the transposition of any two points in the same F-class.

Since Ω is homogeneous in a finite relational language, it is ω-categorical. We may therefore use the group correspondence from the introduction and classify the closed groups H with $G \leq H \leq \mathrm{Sym}(D)$. Here and below closedness is in the topology of pointwise convergence on permutations. For permutation groups on disjoint subsets, we use their product to mean the action on the union that acts independently on each subset; each factor fixes the other subsets pointwise. We shall recover a partition from the transpositions in H. For $x, y \in D$, put

$$x \equiv_H y \quad \iff \quad x = y \ \text{ or the transposition } (x\,y) \text{ belongs to } H.$$

This is an equivalence relation. Only transitivity needs checking: for three distinct points, the identity $(x\,y)(y\,z)(x\,y) = (x\,z)$ proves it. It is H-invariant, because conjugating $(x\,y)$ by $h \in H$ gives $(h(x)\,h(y))$. It contains F, by the preceding observation.

There are only three G-invariant equivalence relations containing F: F, J, and the universal relation. Indeed, if S is joined to one C_i, permutations of the C_i force it to be joined to all of them. If S remains separate but two different C_i are joined, the same permutations force all the C_i to be joined. Otherwise no two F-classes are joined. Apply this observation to $\equiv_H$.

Suppose first that $\equiv_H = F$. Its invariance gives $H \leq G_F$. The stabiliser of the set S in G_F is exactly G: preserving F and S is equivalent to preserving E. Thus, if all elements of H fix S setwise, then $H = G$. Otherwise H acts transitively on the F-classes, since it moves

S to one C_i and contains the permutations of all the C_i. For any $k \in G_F$, choose $h \in H$ with $h(S) = k(S)$. Then $h^{-1}k$ fixes S setwise and belongs to $G \subseteq H$, so $k \in H$. Hence $H = G_F$.

Suppose next that $\equiv_H = J$. Then $H \leq G_J$, and H contains all transpositions within S and within U. Their finite products are dense in $G_U = \mathrm{Sym}(S) \times \mathrm{Sym}(U)$: any prescribed finite restriction of a permutation of either part extends to a permutation of finite support in that part. Closedness therefore gives $G_U \leq H$. The group G_U has index two in G_J, with the other coset consisting of the permutations exchanging S and U. Consequently $H = G_U$ or $H = G_J$.

Finally, if $\equiv_H$ is universal, then H contains every transposition of D, and hence all permutations of finite support. These are dense in $\mathrm{Sym}(D)$, so closedness gives $H = \mathrm{Sym}(D)$. This proves exhaustiveness.

All five listed groups occur as automorphism groups of the displayed definable reducts. To verify distinctness and the asserted incomparabilities, a transposition of one point of C_i with one point of C_j, where $i \neq j$, belongs to G_U but not to G_F. A permutation exchanging the entire F-classes S and C_0 and fixing all other C_i belongs to G_F but not to G_J. A bijection exchanging S with U belongs to G_J but not to G_U. Finally, a transposition of one point of S with one point of U belongs to neither G_F nor G_J. Together with the displayed inclusions, these witnesses prove the claimed lattice. □

7.3 A homogeneous Ramsey expansion

To obtain the Ramsey counterexample, equip S with an order of type $(\mathbb{Q}; <)$, identify U with $\mathbb{Q} \times \mathbb{Q}$ so that E is equality of first coordinates, and give U the lexicographic order. Put every point of S before every point of U, and denote the resulting order on D by $\prec$. Set

$$\Omega^{\prec} = (D; E, U, \prec).$$

The set-theoretic enumeration $C_0, C_1, \ldots$ used above is independent of the rational order of the classes.

Proposition 7.3. *The structure $\Omega^{\prec}$ is homogeneous, Ramsey, and finitely bounded in a finite relational language of maximum arity two. Its reducts include all the structures Γ_n from Theorem 7.1, whose embedding monoids form the strictly increasing chain.*

Proof. For homogeneity, let p be a finite partial isomorphism. Its restriction to S extends to an order automorphism of S. On U, it induces an increasing partial bijection between finitely many E-classes. Extend this to an automorphism of the quotient order, a copy of $(\mathbb{Q}; <)$. Within each matched pair of classes, extend the prescribed increasing partial map to an order isomorphism of their copies of $\mathbb{Q}$; for the remaining classes choose any such isomorphisms. Their union, together with the chosen map on S, is an automorphism extending p.

The age consists of ordered sums $A_0 + A_1$, with the two parts distinguished by U, where A_0 is a finite linear order of E-singletons and A_1 is a finite convexly ordered equivalence relation. A singleton class in A_1 still has its point in U, since it can be a finite piece of an infinite class.

We verify the Ramsey property of these sums. The class of finite linear orders is Ramsey by the finite Ramsey theorem, and the class of finite convexly ordered equivalence relations is Ramsey by [28, Corollary 6.8]. Fix $A = A_0 + A_1$, $B = B_0 + B_1$ in the age and an integer $r \geq 2$. If A does not embed into B, use B as a witness. We use the arrow notation $X \longrightarrow (Y)^Z_q$ to mean that every q-colouring of $\mathrm{Emb}(Z, X)$ admits an embedding of Y on whose image all embeddings of Z have one colour. If A embeds into B, choose

$$D_0 \longrightarrow (B_0)^{A_0}_r, \qquad t = |\,\mathrm{Emb}(A_0, D_0)|, \qquad D_1 \longrightarrow (B_1)^{A_1}_{r^t}.$$

These choices are justified by the two cited Ramsey properties.

An embedding of A into $D_0 + D_1$ is exactly a pair of embeddings of A_0 and A_1 into the corresponding parts. Given an r-colouring of these pairs, colour each embedding of A_1 into D_1 by the vector of its colours over all t embeddings of A_0 into D_0. Choose an embedded $B_1' \subseteq D_1$ on which this vector is constant. On pairs using B_1', the original colour now depends only on the first coordinate. The choice of D_0 gives an embedded $B_0' \subseteq D_0$ on which that remaining colour is constant. Thus $B_0' + B_1'$ is monochromatic. Empty parts cause no difficulty: the empty structure has exactly one embedding into every structure. All structures carry a linear order and are rigid, so the embeddings and copies formulations agree.

For finite boundedness, the age is specified by the universal requirements that $\prec$ is a strict linear order, E is an equivalence relation with convex classes, U is an E-saturated final segment, and every E-class outside U is a singleton. Each violation is witnessed on at most three points. Conversely, every finite structure satisfying these requirements has the described form $A_0 + A_1$ and embeds into $\Omega^{\prec}$. Forbidding the finitely many violating structures on at most three points therefore defines the age.

Finally, all the relations S_n are quantifier-free definable in this expansion. They give the same reducts Γ_n and hence the same strictly increasing chain of self-embedding monoids. □

Appendix B proves that $\Omega^{\prec}$ also has finitely many reducts up to first-order interdefinability. Their exact number is not determined here.

A A direct proof of the Ramsey input

Throughout this appendix, $\prec$ again denotes the convex order on S_2 from Section 3; it is not the order on the counterexample's domain D from Section 7.

Proposition 3.3 uses the Ramsey property of the homogeneous convex-order expansion $\Delta^{\prec} = (S_2; \leq, C, \prec)$. This property is stated in [8, Corollary 3.11], whose proof invokes a general model-companion transfer assertion subsequently disproved in [34, Section 1.3]. We complete the direct join-completion argument of [7, Corollary 3.3] by realising the entire finite Ramsey witness simultaneously, so that arbitrary colourings have a compatible pullback.

The binary refinement and colouring argument follow [30, Lemmas 6.9–6.10 and Theorem 6.11]. Here the marked points may be comparable. Sending them all to leaves, as in the realisation step of [30, Observation 6.5], would lose these comparisons. We instead insert a nonbranching marked vertex immediately above each vertex of the witness, retaining the ancestor relations between marks. The additional checks concern the marked-point realisation and the extension of embeddings to join completions.

A convex order on a finite tree with its root above its descendants is obtained by ordering the child branches at each vertex, listing each entire branch before the next, and listing the vertex after its descendants. We call this the *postorder*. Its restriction to the marked vertices of a normalised diagram is the order used here. Equivalently, it extends $<$, an incomparable point lies outside the interval of a comparable pair, and the outgroup of a C-triple lies outside the interval of the other two entries. The homogeneous expansion $\Delta^{\prec}$ and this description of its finite substructures are given in [8, Section 3.3]. Let $\mathcal{K}^{\prec} = \mathrm{Age}(\Delta^{\prec})$.

We use the following form of Leeb's theorem [31]; see [25, pp. 74–76] for its formulation in terms of finite ordered trees, and [8, Theorem 3.10] for the semilattice formulation. Let $\mathcal{T}$ be the class of all finite rooted trees, with arbitrary finite branching, equipped with the ancestor order (root maximal), the binary least-common-ancestor operation $\vee$, and a postorder. For any $T_A, T_B \in \mathcal{T}$ and integer $r \geq 2$, some $T \in \mathcal{T}$ has the following property: every r-colouring of the embeddings $T_A \to T$ admits an embedding $g : T_B \to T$ for which all composites gh, with $h : T_A \to T_B$ an embedding, have the same colour. Embeddings preserve $\vee$ and the

postorder; they need not send root to root. The ordered trees are rigid, so the formulation in terms of copies also gives this formulation in terms of embeddings. The unrestricted branching of the witness T is important.

Lemma A.1 (Completion of a finite marked tree)**.** *For each nonempty $A \in \mathcal{K}^{\prec}$, there is a finite rooted tree $T_A \in \mathcal{T}$ with at most two children at each vertex, containing the points of A as marked vertices, none of which is a branching vertex. On the marks, restrict the ancestor order and postorder and interpret C by the finite-diagram convention from Section 2; this induced $(\leq, C, \prec)$-structure is precisely A. The whole tree T_A is generated by its marks under $\vee$. For all nonempty $A, B \in \mathcal{K}^{\prec}$, every embedding $\alpha : A \to B$ in the relational signature extends uniquely to an embedding $\widehat{\alpha} : T_A \to T_B$ preserving $\vee$ and the postorder.*

Proof. Choose an embedding of A into $\Delta^{\prec}$, which exists by the definition of its age, and identify A with that finite image. Choose its normalised diagram from Lemma 2.1, retain only the minimal subtree spanning the marks, and suppress unmarked vertices with one child. The order $\prec$ fixes the order of its child branches. Every remaining unmarked vertex is the join of two incomparable marked descendants in its two branches. Thus the resulting tree is generated by its marks.

We check that it and its embeddings are determined by the marked structure, not by a choice of diagram. For a vertex v, write $D_v = \{a \in A : a \leq v\}$ for its nonempty set of marked descendants. If $v = a$ is marked, then D_v is given by the order on A. If $v = x \vee y$ for incomparable $x, y \in A$, then

$$D_v = \{a \in A : \neg C(a; x, y) \ \wedge \ \neg(x < a \wedge y < a)\}.$$

Indeed, a marked point strictly above v satisfies the second excluded condition; a marked point incomparable with v satisfies $C(a; x, y)$; and a marked point below v satisfies neither. A mark cannot equal this branching vertex.

Different vertices have different descendant sets. For two comparable vertices, if the upper vertex is marked, it contributes its own mark; otherwise its branch not containing the lower vertex contains a mark. In either case the upper descendant set is strictly larger. Incomparable vertices have disjoint descendant sets. These sets therefore determine the tree order by inclusion and determine every join as the least vertex above its two arguments. On incomparable vertices the postorder is determined by the order of any marked descendants in their two disjoint branches.

Now take nonempty $A, B \in \mathcal{K}^{\prec}$ and an embedding $\alpha : A \to B$. In T_B, let J be the smallest set containing $\alpha[A]$ and closed under $\vee$. Equip J with the induced ancestor order and draw an edge for each immediate comparison in that order; retain just $\alpha[A]$ as the marked vertices. This is the join-generated tree meant here, rather than the graph subtree containing every original edge on the connecting paths. Every unmarked vertex in J is a join of two incomparable marked descendants, so it is a branching vertex in this tree. For $v \in J$, its set of marked descendants is $D_v \cap \alpha[A]$, where D_v denotes all its marked descendants in T_B. The preceding formulas describe these intersections using the relations on $\alpha[A]$. Since α preserves and reflects $\leq, C, \prec$, they give exactly the descendant sets of the tree constructed for A. This supplies the required embedding, explicitly sending $x \vee y$ to $\alpha(x) \vee \alpha(y)$. It is unique because the marks generate T_A. □

Proposition A.2. *The class $\mathcal{K}^{\prec}$ is Ramsey.*

Proof. Fix nonempty $A, B \in \mathcal{K}^{\prec}$ and an integer $r \geq 2$. If A does not embed into B, take B itself as witness. Otherwise apply Leeb's theorem to T_A, T_B, r, obtaining T. We turn T, whose branching need not be binary, into a marked binary diagram as follows.

The binary refinement uses the method of [30, Lemma 6.9], with joins in place of meets. At each vertex with more than two child branches, replace the branching by any finite binary tree whose leaves carry these branches in their existing order. Keep the old vertex at the top of this refinement. The ancestor relation between old vertices and their postorder remain unchanged. To retain comparable marked points, we now modify the realisation step used for leaves. For each old vertex t other than the root, subdivide the edge from t to its parent in the refined tree by a new vertex t^*, carrying the mark t. For the old root t, add a new parent t^*, carrying that mark, and make t^* the new root. Only the new vertices are marked. They have at most one child, so the diagram is normalised. Its marked structure, with the induced postorder, is some $D \in \mathcal{K}^{\prec}$. To see the ordered realisation explicitly, use the recursive construction in Lemma 2.1: for two child branches choose two incomparable points below the parent and assign them in their $\prec$-order. Convexity orders their entire descendant branches in the same way, and $<\subseteq\prec$ orders each vertex after its descendants. Thus the construction realises the prescribed postorder as well as $\leq, C$.

For every embedding $h : T_A \to T$, define $h^* : A \to D$ by $h^*(a) = h(a)^*$. This is an embedding of the marked structures. Order, incomparability and postorder are unchanged by refinement and insertion. If $C(x; y, z)$ holds in A, the vertex $y \vee z$ in T_A is incomparable with x. Its image has the same property in T. After refinement, the join of $h(y)^*, h(z)^*$ lies at or below the old vertex $h(y \vee z)$, still incomparable with $h(x)^*$. Hence the same C-orientation holds in D. On an antichain triple both structures have exactly one true orientation, so preservation implies reflection. On other triples C is false in both. This proves the claim about h^*. The same argument applies to embeddings $T_B \to T$.

The remaining colouring argument is the same composition-compatible pullback as in [30, Lemma 6.10 and Theorem 6.11]. Given a colouring $\chi : \mathrm{Emb}(A, D) \to \{1, \dots, r\}$, colour each $h : T_A \to T$ by $\chi(h^*)$. Choose the monochromatic embedding $g : T_B \to T$ supplied by Leeb's theorem. For any $\alpha : A \to B$, Lemma A.1 gives $\widehat{\alpha} : T_A \to T_B$, and on each $a \in A$ we have

$$(g\widehat{\alpha})^*(a) = g(\alpha(a))^* = (g^*\alpha)(a).$$

Consequently all embeddings $g^*\alpha$ have the same χ-colour. Since $g^* : B \to D$ is an embedding, D is the required Ramsey witness. The empty structure, if allowed, gives only the unique empty embedding and requires no separate construction. □

B First-order finiteness for the ordered example

Recall the ordered structure $\Omega^{\prec}$ from Section 7. Its domain is $D = S \sqcup U$, where S is a copy of the rational order and U is the lexicographically ordered union of the rationally indexed classes C_q, each itself a copy of the rational order; all of S precedes U. We call S and the individual C_q the *blocks*. These are the classes of F, not of E: the E-classes in S are singletons.

Theorem B.1. *The homogeneous Ramsey structure $\Omega^{\prec}$ has finitely many reducts up to first-order interdefinability.*

B.1 Strategy and classification inputs

Proof strategy. We must control every closed group G containing $\mathrm{Aut}(\Omega^{\prec})$. Each original block admits arbitrary increasing automorphisms fixing all other blocks pointwise. We first examine how G can mix these blocks: join points lying together in an image $g[B]$, with $g \in G$, of an original block, and take connected components. The internal motions on these images let us transport any finite subset of a component to any other subset of the same size. The

original symmetries force just three component partitions: the original blocks, the two parts S, U, or the whole domain.

Cameron's classification controls the actions inside these components. When the original blocks remain separate, their independent internal motions form a normal subgroup of G. After factoring it out, the remaining action records block permutations and at most two possible orientation states per block. We represent those states by one- or two-element fibres over the block indices and apply Bodor's bound for finite covers. The distinguished block S requires separate treatment according to whether its index is fixed or moved. Finally, the uniform bound on tuple orbits and a fixed subgroup with finitely many orbits in each arity give finitely many actual overgroups, rather than merely finitely many groups up to relabelling.

The contribution of this appendix is this reduction for $\Omega^{\prec}$, including the component argument and the treatment of the distinguished block. The general classifications are external inputs: Cameron supplies the possible component groups, and Bodor supplies a uniform bound for the resulting finite covers. We state their precise consequences next.

A permutation group on an infinite set is *highly set-transitive* if it is transitive on the n-element subsets for every positive integer n. Cameron's theorem [18] says that the closed highly set-transitive groups acting on a countably infinite set are, up to relabelling,

$$\mathrm{Aut}(\mathbb{Q};<), \quad \mathrm{Aut}(\mathbb{Q};\mathrm{Betw}), \quad \mathrm{Aut}(\mathbb{Q};\mathrm{Cyc}), \quad \mathrm{Aut}(\mathbb{Q};\mathrm{Sep}), \quad \mathrm{Sym}(\mathbb{Q}).$$

Here $\mathrm{Betw}(x,y,z)$ means $x<y<z$ or $z<y<x$; $\mathrm{Cyc}(x,y,z)$ means $x<y<z$, $y<z<x$, or $z<x<y$; and $\mathrm{Sep}(x,y,u,v)$ means that the four points are distinct and exactly one of $\mathrm{Cyc}(x,u,y)$, $\mathrm{Cyc}(x,v,y)$ holds. The same theorem classifies the closed overgroups of the given action $\mathrm{Aut}(\mathbb{Q};<)$: they are exactly the five displayed groups for that order.

A group is *oligomorphic* if it has finitely many orbits on tuples of each fixed length. Two consequences of Cameron's theorem will be useful. First, each group is the automorphism group of a structure with relations of arity at most four. Thus only finitely many closed highly set-transitive groups can contain a fixed oligomorphic group P: the defining relations are unions of the finitely many P-orbits in arities at most four. Second, for each of the five groups K, its normalizer in the full symmetric group satisfies

$$[N_{\mathrm{Sym}(\mathbb{Q})}(K) : K] \leq 2. \tag{5}$$

For the order group, the two orbits on injective pairs record the two order orientations; for the circular-order group, the two orbits on injective triples record the two cyclic orientations. A normalizing permutation preserves or exchanges the two relevant orbits, giving respectively the betweenness and separation group as normalizer. The betweenness group is self-normalizing: its three orbits on injective triples are labelled by the coordinate occupied by the middle point. Coordinate permutations induce the full symmetric group on these three labels. A normalizer's permutation of the orbit labels commutes with this action, whose centralizer in the symmetric group on the labels is trivial. Thus the normalizer preserves each orbit, including the betweenness relation. For the separation group, its three orbits on injective quadruples are labelled by the three opposite-pair partitions of the coordinate positions. Coordinate permutations again induce the full symmetric group on those three labels, so the same argument proves self-normalization. The full symmetric group is self-normalizing.

A *finite cover* of a permutation group L on I is a permutation group V on X, together with a surjection $\pi : X \to I$ with finite fibres, such that V preserves the fibre equivalence relation and induces exactly L on I. Transitivity on X is not required. The groups arising below act on these finite auxiliary fibres; the original blocks S, C_q are infinite and are not themselves such fibres.

Finite covers of the highly set-transitive groups were classified by Ivanov [26, Section 2]. We use the group-theoretic formulation in [17, Theorem 5.2], together with Bodor's tuple-orbit bound [17, Lemma 5.29], in the following form.

Proposition B.2 (Finite-cover bound)**.** *Let V act on X as a closed finite cover of a closed highly set-transitive group on a countably infinite set, with fibres of size at most two. For $\bar{x}, \bar{y} \in X^n$, the tuples lie in the same V-orbit if and only if, for every $1 \le k \le \min\{16, n\}$ and every choice $1 \le i_1 < \cdots < i_k \le n$, the tuples $(x_{i_1}, \ldots, x_{i_k})$ and $(y_{i_1}, \ldots, y_{i_k})$ lie in the same V-orbit. The group element witnessing this may depend on the chosen positions.*

Here is the application of the two cited results. Transitivity of the induced group makes all fibres equicardinal. Theorem 5.2 of [17] identifies them with a fixed finite set A, identifies the base with $\mathbb{Q}$, and puts V in a normal form on $A \times \mathbb{Q}$. In that form the kernel of the action on the base is

$$\{(\ell_q)_{q\in\mathbb{Q}} \in L_0^{\mathbb{Q}} : \ell_q H_0 = \ell_r H_0 \text{ for all } q, r \in \mathbb{Q}\}, \qquad H_0 \trianglelefteq L_0 \le \mathrm{Sym}(A).$$

The sequence $(\ell_q)_{q\in\mathbb{Q}}$ denotes the permutation $(a, q) \mapsto (\ell_q(a), q)$ of $A \times \mathbb{Q}$. Here L_0 is the action of the kernel on one fibre; H_0 consists of the restrictions to that fibre of elements of V fixing every point outside it. Neither group is required to be transitive, so the theorem includes intransitive covering groups.

Lemma 5.29 applies to the full group V in this normal form: its tuple orbits are determined by corresponding coordinate selections of length $4|A|[L_0 : H_0]$, allowing repetitions. Since $|A| \le 2$ and $[L_0 : H_0] \le |\mathrm{Sym}(A)| \le 2$, this length is at most 16. Relabelling does not change the conclusion. Repeated selections of a coordinate position require no additional checks: an element matching the entries at the distinct selected positions also matches their repeated occurrences. This concerns repetition of positions in a selection; equal entries at different positions of the original tuples are allowed throughout.

We will use the bound through an elementary counting argument. Suppose that all the possible V act on the same set X and contain a fixed oligomorphic group P_0. In each arity at most 16, every V-orbit is a union of P_0-orbits. The collections of P_0-orbits contained in individual V-orbits therefore partition the finite set of P_0-orbits, so only finitely many such partitions are possible. Proposition B.2 then determines the orbits in every arity. A closed permutation group is determined by all its finite-tuple orbits: a permutation preserving every such orbit agrees on each finite set with a group element and hence belongs to the group by closedness. Thus only finitely many actual groups V can occur.

B.2 Combining motions on overlapping sets

Say that a permutation is *supported on* B if it fixes every point outside B. The following elementary lemma explains why the components in the strategy above have highly set-transitive supported actions.

Lemma B.3. *Let $G \le \mathrm{Sym}(X)$, and let $\mathcal{A}$ be a family of infinite subsets of X. Suppose that each $A \in \mathcal{A}$ supports a subgroup of G acting highly set-transitively on A. Form the graph on X in which distinct points are adjacent if they lie together in a member of $\mathcal{A}$. For every connected component Y meeting a member of $\mathcal{A}$, the subgroup of G supported on Y acts highly set-transitively on Y.*

Proof. If a finite subset $T \subseteq Y$ contains x but not y, and $x, y \in A \in \mathcal{A}$, choose an element g of the subgroup of G supported on A that sends $T \cap A$ to $((T \cap A) \setminus \{x\}) \cup \{y\}$. The assumed highly set-transitive action supplies such an element. Since g fixes every point outside A, we have

$$g[T] = (T \setminus \{x\}) \cup \{y\}.$$

Only this equality of sets is required; the other members of $T \cap A$ need not be fixed by g. Such replacements connect any two n-element subsets of a connected graph. To check this, restrict to a finite connected subgraph containing both subsets and choose a spanning tree. Induct on the number of vertices of this tree. At a leaf where the two subsets agree about occupancy, keep that occupancy unchanged and remove the leaf. If only the first subset contains the leaf, choose a nearest vacant vertex, meaning a vertex outside the current finite subset; all preceding vertices on the path are occupied. Moving points successively towards that vacancy makes the leaf vacant without collisions, and induction applies after removing it. If only the second subset contains the leaf, reverse the roles of the subsets and then reverse the resulting sequence of moves. This proves the lemma, including the empty and full finite configurations. □

B.3 Proof of finiteness

Proof of Theorem B.1. Put $H = \mathrm{Aut}(\Omega^{\prec})$, and let G be a closed group with $H \leq G \leq \mathrm{Sym}(D)$. Write $\mathcal{B} = \{S\} \cup \{C_q : q \in \mathbb{Q}\}$ for the set of F-classes. Form the graph whose edges join points in a common set $g[B]$, for $g \in G$ and $B \in \mathcal{B}$, and denote its component equivalence by T_G. This equivalence is G-invariant and contains F. Each original block supports a copy of $\mathrm{Aut}(\mathbb{Q}; <)$ inside H; conjugation supplies such an action on every $g[B]$. Lemma B.3 therefore applies to every component.

The only H-invariant equivalences containing F are F, J, D^2. To see this, the action on the F-classes fixes S and is the rational-order automorphism group on the other classes. It is transitive on these other classes and on their unordered pairs. Joining S to one class therefore joins it to all classes; joining two other classes joins all of U. Thus there are three cases.

Case 1: one component. If $T_G = D^2$, Lemma B.3 makes G highly set-transitive. The finite-orbit consequence of Cameron's theorem above gives only finitely many possibilities, since H is oligomorphic.

Case 2: the components are S, U. Suppose $T_G = J$. For $Y = S, U$, let K_Y be the group on Y induced by the elements of G supported on Y. Extension by the identity identifies K_Y homeomorphically with

$$G \cap \{g \in \mathrm{Sym}(D) : g(x) = x \text{ for every } x \in D \setminus Y\}.$$

This intersection is closed, so K_Y is closed in $\mathrm{Sym}(Y)$. It is highly set-transitive by Lemma B.3. Write $H|_Y = \{h|_Y : h \in H\}$ for the induced action of H on Y. The actions of H on S and U are independent, so every restriction in $H|_Y$ extends by the identity outside Y to an element of H. Thus $H|_Y \leq K_Y$. Since $H|_Y$ is oligomorphic, Cameron's theorem gives finitely many possible pairs (K_S, K_U). The product $K_S \times K_U$ lies in G, and conjugation by G preserves it, possibly exchanging the factors. Its orbits on D are exactly S, U, so its full normalizer also preserves that unordered partition. The normalizer's action on the two components has image of size at most two. Its subgroup preserving both components setwise, modulo $K_S \times K_U$, is

$$(N_{\mathrm{Sym}(S)}(K_S)/K_S) \times (N_{\mathrm{Sym}(U)}(K_U)/K_U).$$

Each factor has size at most two by (5). Allowing the possible exchange of the components, the full quotient has size at most $2 \cdot 2 \cdot 2 = 8$. For each fixed pair there are therefore only finitely many intermediate groups G.

Case 3: the original blocks remain the components. It remains to treat $T_G = F$. For every $B \in \mathcal{B}$, let K_B be the supported group on B. It is closed by the extension-by-identity argument in Case 2. It contains the original rational-order automorphism group on B, so it is one of Cameron's five groups for that given order. The choice on the C_q is uniform, since

H transports them by increasing isomorphisms. There are thus at most 25 local assignments, allowing a separate choice on S. Fix one. Closedness gives

$$K := \prod_{B \in \mathcal{B}} K_B \ \leq G :$$

products supported on finitely many blocks already lie in G, and approximate every element of the displayed product on each finite subset of D.

The orbits of K are the blocks $\mathcal{B}$. Conjugation by G transports their supported groups, so G lies in the block-permuting normalizer W of K. Identify their index set with

$$I = \{p\} \sqcup \mathbb{Q},$$

where p denotes S. We now construct the finite fibres that record what remains after quotienting by K. Write $B_p = S$ and $B_q = C_q$, and fix an increasing identification $\theta_i : \mathbb{Q} \to B_i$ for each $i \in I$. Put

$$K_i = \theta_i^{-1} K_{B_i} \theta_i, \qquad N_i = N_{\mathrm{Sym}(\mathbb{Q})}(K_i), \qquad Q_i = N_i / K_i, \qquad X = \bigsqcup_{i \in I} (\{i\} \times Q_i).$$

The fibre over i is thus a copy of the quotient group Q_i, of size at most two by (5). Cameron's five groups are pairwise nonconjugate: in the order listed above, their numbers of orbits on injective pairs and triples are $(2, 6), (1, 3), (1, 2), (1, 1), (1, 1)$, and the last two have respectively three and one orbits on injective quadruples. Therefore a member $w \in W$ can send B_i to B_j only if $K_i = K_j$. Let σ_w be its permutation of I. We then have $K_i = K_{\sigma_w(i)}$, and hence $N_i = N_{\sigma_w(i)}$ and $Q_i = Q_{\sigma_w(i)}$ as actual groups in these coordinates. Its restriction from B_i to $B_{\sigma_w(i)}$, written in the fixed coordinates, is

$$n_i(w) = \theta_{\sigma_w(i)}^{-1} w \theta_i \in N_i.$$

Define its action on the auxiliary fibres by

$$\Phi(w)(i, a) = (\sigma_w(i), [n_i(w)]a) \qquad (i \in I,\ a \in Q_i),$$

where brackets denote the coset modulo K_i. Composition of the local maps gives a homomorphism $\Phi : W \to \mathrm{Sym}(X)$ with kernel K.

Colour every point of the fibre $\{i\} \times Q_i$ by the group K_i, one of the five groups in Cameron's list. The image $Z = \Phi(W)$ consists exactly of the permutations of X preserving the fibre equivalence and these colours. To see surjectivity onto this group, choose the prescribed permutation of blocks with equal colours. Every prescribed bijection between their auxiliary fibres is multiplication by an element of their common group Q_i, since $|Q_i| \leq 2$. Choose a representative in N_i independently for each block and combine the resulting block bijections. Their union lies in W and induces the prescribed permutation of X. Thus Z is closed: failure to preserve the fibre equivalence or a colour is witnessed on finitely many points. If the group chosen for S differs from the common group chosen for the C_q, the fibre over p is fixed setwise.

For the subsequent closed-group arguments we need $\Phi : W \to Z$ to be continuous and open. Each K_i is closed of finite index in N_i, hence open. Choose a finite nonempty $E_i \subseteq \mathbb{Q}$ such that every element of N_i fixing E_i pointwise belongs to K_i. For any finite set of indices $J \subseteq I$, fixing $\bigcup_{i \in J} \theta_i[E_i]$ pointwise forces $\Phi(w)$ to fix all auxiliary points over J. This proves continuity at the identity, and hence everywhere. Conversely, let $A \subseteq D$ be finite and let J be the indices of its blocks. Every element of Z fixing the fibres over J pointwise has a representative in W equal to the identity on those entire blocks. Hence the image of the

pointwise stabilizer of A contains an open pointwise stabilizer in Z. That image is a subgroup and is therefore open; translating such neighbourhoods proves openness of Φ.

Set $V = \Phi(G)$. Since $K \leq G$, we have $G = \Phi^{-1}(V)$, and

$$Z \setminus V = \Phi(W \setminus G).$$

The right-hand side is open, so V is closed in Z, hence in $\mathrm{Sym}(X)$. It remains to show that only finitely many such V occur.

The image of H is the same subgroup P_0 for all these choices of G with the fixed local assignment: it fixes the fibre over p pointwise and acts on the remaining fibres by

$$(q, a) \longmapsto (h(q), a) \qquad (h \in \mathrm{Aut}(\mathbb{Q}; <)).$$

Indeed, every local map of an element of H, expressed in the coordinates θ_i, is increasing and belongs to K_i; its coset in Q_i is the identity. The tuple orbits of P_0 are determined by the finite fibre coordinates, the occurrences over p, and the equality and order pattern of the rational indices. Thus P_0 is oligomorphic and is contained in every V.

Pointwise-convergence topologies on permutation groups of countable sets are metrizable, so closedness can be checked by taking limits of convergent sequences. We use this both here and for the projections below. The induced group L of V on I is closed as well. Suppose $l_n \in L$ converges to $l \in \mathrm{Sym}(I)$. Choose $v_n \in V$ inducing l_n, and write

$$v_n(i, a) = (l_n(i), c_{n,i}a), \qquad c_{n,i} \in Q_i.$$

As each Q_i is finite and I is countable, a diagonal subsequence makes $c_{n,i}$ eventually constant for every i. Write its eventual value as c_i. The limit l preserves the fibre colours, as each l_n does. The maps therefore converge pointwise to the permutation $(i, a) \mapsto (l(i), c_i a)$ of X. Closedness of V puts this permutation in V, so its induced permutation l belongs to L.

The distinguished index moves. If L moves p, it is highly set-transitive. For every $n \geq 1$, its subgroup $\mathrm{Aut}(\mathbb{Q}; <)$ fixing p has exactly two orbits on n-element subsets: those containing p and those avoiding it. Choose $l \in L$ with $l(p) \neq p$, and an n-element set containing p but avoiding $l^{-1}(p)$. Its image avoids p, joining the two orbits. In this case all fibres have the same size, at most two, and V is a closed finite cover of L. Proposition B.2 and the counting argument following it, with the fixed subgroup P_0, give only finitely many possibilities for V.

The distinguished index is fixed. Finally suppose L fixes p. Let $A_p = \mathrm{Sym}(\{p\} \times Q_p)$ be the full symmetric group of the fixed fibre over p, so $|A_p| \leq 2$, and write $X_{\mathbb{Q}}$ for the union of the remaining fibres. Thus $V \leq A_p \times \mathrm{Sym}(X_{\mathbb{Q}})$. Let $V_{\mathbb{Q}}$ be the projection of V to the second factor, and put

$$N = \{v \in V_{\mathbb{Q}} : (1, v) \in V\}.$$

Both groups are closed: for $V_{\mathbb{Q}}$, a convergent sequence of projections has lifts with a constant first coordinate after taking a subsequence; for N, extend by the identity on the finite fibre. Since P_0 fixes the fibre over p pointwise, its restriction to $X_{\mathbb{Q}}$ belongs to N, and hence also to $V_{\mathbb{Q}}$. Thus both groups contain the action $(q, a) \mapsto (h(q), a)$ for $h \in \mathrm{Aut}(\mathbb{Q}; <)$. Their actions on the rational indices are closed by the finite-fibre argument just given and contain $\mathrm{Aut}(\mathbb{Q}; <)$, so Cameron's theorem makes those actions highly set-transitive. Thus $V_{\mathbb{Q}}$ and N are closed finite covers with fibres of size at most two. Proposition B.2, with the common diagonal subgroup just identified, gives finitely many actual possibilities for each.

Moreover, $N \trianglelefteq V_{\mathbb{Q}}$ and $[V_{\mathbb{Q}} : N] \leq |A_p| \leq 2$: two lifts having the same first coordinate have second coordinates in the same coset of N. For a fixed pair $(V_{\mathbb{Q}}, N)$, consider the quotient homomorphism

$$\rho : A_p \times V_{\mathbb{Q}} \longrightarrow A_p \times (V_{\mathbb{Q}}/N), \qquad \rho(a, v) = (a, vN).$$

Its kernel is $\{1\} \times N$, contained in V by the definition of N. Hence $V = \rho^{-1}(\rho(V))$. The target is a finite group, so its finitely many subgroups give only finitely many possibilities for V. This proves that only finitely many closed overgroups G occur, and hence proves the theorem. □

Acknowledgment

The author thanks Manuel Bodirsky and Michael Pinsker for valuable comments and suggestions on the original manuscript. The original work was supported by the EFOP-3.6.2-16-2017-00015 project, co-financed by the European Social Fund; by the Hungarian National Research, Development and Innovation Fund under grants FK 124814 and PD 125160; by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences; and by the ÚNKP-19-4 New National Excellence Program.

The author was supported by the NRDI grant KKP 138270.

References

[1] S. A. Adeleke and D. Macpherson, *Classification of infinite primitive Jordan permutation groups*, Proc. London Math. Soc. (3) **72** (1996), no. 1, 63–123. doi:10.1112/plms/s3-72.1.63.

[2] S. A. Adeleke and P. M. Neumann, *Relations related to betweenness: their structure and automorphisms*, Mem. Amer. Math. Soc. **131** (1998), no. 623, viii+125 pp. doi:10.1090/memo/0623.

[3] L. Agarwal, *Reducts of the generic digraph*, Ann. Pure Appl. Logic **167** (2016), 370–391. doi:10.1016/j.apal.2015.12.006.

[4] L. Agarwal and M. Kompatscher, $2^{\aleph_0}$ *pairwise nonisomorphic maximal-closed subgroups of* $\mathrm{Sym}(\mathbb{N})$ *via the classification of the reducts of the Henson digraphs*, J. Symbolic Logic **83** (2018), no. 2, 395–415. doi:10.1017/jsl.2017.74.

[5] M. Bodirsky, *Cores of countably categorical structures*, Log. Methods Comput. Sci. **3** (2007), no. 1, Article 2, 1–16. doi:10.2168/LMCS-3(1:2)2007.

[6] M. Bodirsky, *Complexity of Infinite-Domain Constraint Satisfaction*, Lecture Notes in Logic **52**, Cambridge University Press, 2021. doi:10.1017/9781107337534.

[7] M. Bodirsky, D. Bradley-Williams, M. Pinsker and A. Pongrácz, *The universal homogeneous binary tree*, preprint, 7 September 2014, Corollary 3.3. arXiv:1409.2170v1. The cited completion argument is in this early version, rather than the journal version [8].

[8] M. Bodirsky, D. Bradley-Williams, M. Pinsker and A. Pongrácz, *The universal homogeneous binary tree*, J. Logic Comput. **28** (2018), 133–163. doi:10.1093/logcom/exx043. Preprint: arXiv:1409.2170.

[9] M. Bodirsky, P. Jonsson and T. V. Pham, *The reducts of the homogeneous binary branching C-relation*, J. Symbolic Logic **81** (2016), 1255–1297. doi:10.1017/jsl.2016.37. Numbered preprint references are to arXiv:1408.2554v2.

[10] M. Bodirsky, P. Jonsson and T. V. Pham, *The complexity of phylogeny constraint satisfaction problems*, ACM Trans. Comput. Logic **18** (2017), no. 3, Article 23, 23:1–23:42. doi:10.1145/3105907.

[11] M. Bodirsky and M. Junker, $\aleph_0$-*categorical structures: endomorphisms and interpretations*, Algebra Universalis **64** (2010), 403–417. doi:10.1007/s00012-011-0110-y.

[12] M. Bodirsky and J. Kára, *The complexity of temporal constraint satisfaction problems*, J. ACM **57** (2010), no. 2, Article 9, 9:1–9:41. doi:10.1145/1667053.1667058.

[13] M. Bodirsky and J. Nešetřil, *Constraint satisfaction with countable homogeneous templates*, J. Logic Comput. **16** (2006), no. 3, 359–373. doi:10.1093/logcom/exi083.

[14] M. Bodirsky and M. Pinsker, *Minimal functions on the random graph*, Israel J. Math. **200** (2014), 251–296. doi:10.1007/s11856-014-1042-y.

[15] M. Bodirsky, M. Pinsker and A. Pongrácz, *The 42 reducts of the random ordered graph*, Proc. London Math. Soc. (3) **111** (2015), 591–632. doi:10.1112/plms/pdv037.

[16] B. Bodor, *Classification of ω-categorical monadically stable structures*, J. Symbolic Logic **89** (2024), 460–495. doi:10.1017/jsl.2023.66.

[17] B. Bodor, *Structures with not too fast unlabelled growth*, J. Symbolic Logic, First View (2026), 1–52. doi:10.1017/jsl.2026.10239.

[18] P. J. Cameron, *Transitivity of permutation groups on unordered sets*, Math. Z. **148** (1976), 127–139. doi:10.1007/BF01214702.

[19] P. J. Cameron, *Some treelike objects*, Quart. J. Math. Oxford (2) **38** (1987), 155–183. doi:10.1093/qmath/38.2.155.

[20] M. Droste, *Structure of partially ordered sets with transitive automorphism groups*, Mem. Amer. Math. Soc. **57** (1985), no. 334, vi+100 pp. doi:10.1090/memo/0334.

[21] M. Droste, *Partially ordered sets with transitive automorphism groups*, Proc. London Math. Soc. (3) **54** (1987), no. 3, 517–543. doi:10.1112/plms/s3-54.3.517.

[22] M. Droste, W. C. Holland and H. D. Macpherson, *Automorphism groups of infinite semilinear orders (I)*, Proc. London Math. Soc. (3) **58** (1989), no. 3, 454–478. doi:10.1112/plms/s3-58.3.454.

[23] M. Droste, W. C. Holland and H. D. Macpherson, *Automorphism groups of infinite semilinear orders (II)*, Proc. London Math. Soc. (3) **58** (1989), no. 3, 479–494. doi:10.1112/plms/s3-58.3.479.

[24] M. Droste, W. C. Holland and H. D. Macpherson, *Automorphism groups of homogeneous semilinear orders: normal subgroups and commutators*, Canad. J. Math. **43** (1991), no. 4, 721–737. doi:10.4153/CJM-1991-041-7.

[25] R. L. Graham and B. L. Rothschild, *Some recent developments in Ramsey theory*, in M. Hall, Jr. and J. H. van Lint (eds.), *Combinatorics, Part 2: Graph Theory; Foundations, Partitions and Combinatorial Geometry*, Mathematical Centre Tracts **56**, Mathematisch Centrum, Amsterdam, 1974, pp. 61–76. Digitised revised edition (1975).

[26] A. A. Ivanov, *Finite covers, cohomology and homogeneous structures*, Proc. London Math. Soc. (3) **78** (1999), 1–28. doi:10.1112/S002461159900163X.

[27] M. Junker and M. Ziegler, *The 116 reducts of* $(\mathbb{Q}, <, a)$, J. Symbolic Logic **73** (2008), 861–884. doi:10.2178/jsl/1230396752.

[28] A. S. Kechris, V. G. Pestov and S. Todorcevic, *Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups*, Geom. Funct. Anal. **15** (2005), 106–189. doi:10.1007/s00039-005-0503-1.

[29] M. Kompatscher and T. V. Pham, *A complexity dichotomy for poset constraint satisfaction*, in *34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)*, Leibniz International Proceedings in Informatics **66**, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017, 47:1–47:12. doi:10.4230/LIPIcs.STACS.2017.47.

[30] H. Lee, *Ramsey properties of trees and C-relations via the partite construction*, preprint, 2026. arXiv:2610.18402v1.

[31] K. Leeb, *Vorlesungen über Pascaltheorie*, Arbeitsberichte des Instituts für Mathematische Maschinen und Datenverarbeitung **6**, Friedrich-Alexander-Universität Erlangen–Nürnberg, 1973. The ordered-tree theorem is also stated in [25, pp. 74–76].

[32] J. Linman and M. Pinsker, *Permutations on the random permutation*, Electron. J. Combin. **22** (2015), no. 2, Paper P2.54. doi:10.37236/4910.

[33] D. Macpherson, *A survey of Jordan groups*, in R. Kaye and D. Macpherson (eds.), *Automorphisms of First-Order Structures*, Oxford University Press, Oxford, 1994, pp. 73–110. doi:10.1093/oso/9780198534686.003.0003.

[34] A. Mottet and M. Pinsker, *Cores over Ramsey structures*, J. Symbolic Logic **86** (2021), 352–361. doi:10.1017/jsl.2021.6.

[35] P. P. Pach, M. Pinsker, G. Pluhár, A. Pongrácz and Cs. Szabó, *Reducts of the random partial order*, Adv. Math. **267** (2014), 94–120. doi:10.1016/j.aim.2014.08.008.

[36] A. Pongrácz, *Reducts of the Henson graphs with a constant*, Ann. Pure Appl. Logic **168** (2017), 1472–1489. doi:10.1016/j.apal.2017.01.009.

[37] A. Pongrácz, *Existential reducts of the binary branching semilinear order and the Thomas conjecture*, unpublished manuscript, 2020. Author's manuscript.

[38] A. Pongrácz, *The existential Thomas Conjecture*, in *Homogeneous Structures: Model Theory meets Universal Algebra*, Oberwolfach Reports **18** (2021), no. 1, 22–23. doi:10.4171/OWR/2021/1.

[39] S. Thomas, *Reducts of the random graph*, J. Symbolic Logic **56** (1991), 176–181. doi:10.2307/2274912.

[40] S. Thomas, *Reducts of random hypergraphs*, Ann. Pure Appl. Logic **80** (1996), 165–193. doi:10.1016/0168-0072(95)00061-5.