University of Michigan, Summer 2022 REU Report

Miles Kretschmer
Mentor: Matthew Harrison-Trainor

July 2022

Contents

1 Introduction

In this document, we will present several results in computable model theory and infinitary logic. In this section, we will review the major concepts from model theory, computability theory, and infinitary logic that we will use. In Section 2, we will focus on variations on the computable model theoretic property of relative decidability. We will construct examples which separate various notions of relative decidability for theories and structures, and give a characterization of the property of uniform relative decidability in n Turing jumps, in terms of quantifier elimination. This will require a result in infinitary model theory. In Section 3, we will prove a number of special cases of this, which will introduce necessary ideas for the proof of the general case. We will also present partial results from two unsuccessful approaches to the general case, in digressions. Having introduced the necessary ideas, we will define a notion of forcing, which we will then use to prove the general case.

1.1 Model Theory

In this section, we briefly go over the notion of a diagram in model theory, which will be of importance throughout later sections.

For $\mathcal{A}$ an $\mathcal{L}$-structure, we define the diagram of $\mathcal{A}$, denoted $D(\mathcal{A})$, as follows. Let $\mathcal{L}_{\mathcal{A}}$ be $\mathcal{L}$ together with new constants naming the elements of $\mathcal{A}$. $\mathcal{A}$ is then a $\mathcal{L}_{\mathcal{A}}$-structure in a natural way. We let $D(\mathcal{A})$ be the set of $\mathcal{L}_{\mathcal{A}}$ sentences true in $\mathcal{A}$. This encodes the structure of $\mathcal{A}$, in the following way. If $\mathcal{B} \models D(\mathcal{A})$, we can define an embedding $\mathcal{A} \hookrightarrow \mathcal{B}$ by mapping each element of $\mathcal{A}$ to the interpretation of the corresponding constant.

The elementary diagram of $\mathcal{A}$, denoted $E(\mathcal{A})$, is the theory of $\mathcal{A}$ in the language $\mathcal{L}_{\mathcal{A}}$. This encodes substantially more information about $\mathcal{A}$. For instance, the embedding defined above is an elementary embedding, if $\mathcal{B} \models E(\mathcal{A})$. When working with diagrams, we will habitually identify models of $D(\mathcal{A})$ and $E(\mathcal{A})$ with superstructures and elementary extensions of $\mathcal{A}$ respectively. In this way, the embeddings defined above will be construed as inclusions.

1.2 Computable Model Theory

In computable model theory, we study the computability theoretic properties of countable structures in countable languages. We will identify countable structures with ω-presentations, i.e, with structures whose domain is the natural numbers, ω. Using Gödel numbering, we can assign numbers to formulas in a countable language. Using this, we can for a structure $\mathcal{A}$ identify $D(\mathcal{A})$ and $E(\mathcal{A})$ with sets of natural numbers, as follows. We identify $D(\mathcal{A})$ with the set of codes for pairs (n, m) where φ_n is atomic, m codes a tuple $\overline{k}$ and $\mathcal{A} \models \varphi(\overline{k})$. We identify $E(\mathcal{A})$ with the set of codes for pairs (n, m) where φ_n is any formula, m codes a tuple $\overline{k}$ and $\mathcal{A} \models \varphi(\overline{k})$.

In this formulation, $D(\mathcal{A})$ encodes the functions, constants, and relations of $\mathcal{A}$ as a set of natural numbers, so computations relative to this set can be thought of as functions that make use of the constants, functions, and relations of $\mathcal{A}$. We will now go over a few definitions of computable model theoretic properties which consider. For a more detailed introduction, see [CMS21].

We say that a structure $\mathcal{A}$ is computable if $D(\mathcal{A})$ is computable. That is, if the interpretation of the language in $\mathcal{A}$ is computable. We say that $\mathcal{A}$ is decidable if $E(\mathcal{A})$ is computable. As it stands, decidability implies computability. We say that a structure $\mathcal{A}$ is relatively decidable if $E(\mathcal{A}) \leq_T D(\mathcal{A})$. That is, if we can decide formulas in $\mathcal{A}$ using its functions and relations. We will want to get at truly structural properties of structures, that is, properties of isomorphism types of structure. One such property is if we can decide formulas in a uniform way in isomorphic copies of a structure. We say that a structure is uniformly relatively decidable if there is a Turing functional Φ such that for any $\mathcal{B} \simeq \mathcal{A}$, $\Phi^{D(\mathcal{B})}$ computes $E(\mathcal{B})$.

We can also define associated properties of theories, in terms of their countable models. We say that a theory T is relatively decidable if for every countable model $\mathcal{A} \models T$, $E(\mathcal{A}) \leq_T D(\mathcal{A})$. That is, all of its models are relatively decidable. We say that T is uniformly relatively decidable if there is a Turing functional Φ such for any countable $\mathcal{A} \models T$, $\Phi^{D(\mathcal{A})}$ computes $E(\mathcal{A})$. Note that this is stronger than requiring that every model of T us uniformly relatively decidable, as we want the effective procedure we use to be uniform in the theory.

1.3 Infinitary Languages

In Section 2, we will characterize certain relative decidability properties using quantifier elimination properties. As an intermediate stage in these characterizations, we will use formulas in an infinitary language. In Section 3, we will prove a result concerning the relationship between finitary and infinitary languages. In this section, we introduce these infinitary languages, and the notions of complexity of formulas that our theorems will use.

Let κ be an infinite cardinal. We define the language $\mathcal{L}_{\kappa,\omega}$ by the following recursive clauses.

1. If ψ is an atomic formula of $\mathcal{L}_{\omega,\omega}$, $\psi \in \mathcal{L}_{\kappa,\omega}$.

2. If $\phi \in \mathcal{L}_{\kappa,\omega}$, then $\neg\phi \in \mathcal{L}_{\kappa,\omega}$

3. If $\phi(\overline{y}) \in \mathcal{L}_{\kappa,\omega}$, then $\forall\overline{y}\phi(\overline{y}) \in \mathcal{L}_{\kappa,\omega}$ and $\exists\overline{y}\phi(\overline{y}) \in \mathcal{L}_{\kappa,\omega}$.

4. If $\Phi \subset \mathcal{L}_{\kappa,\omega}$ and $|\Phi| < \kappa$, then $\bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi \in \mathcal{L}_{\kappa,\omega}$ and $\bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi \in \mathcal{L}_{\kappa,\omega}$

In other words, $\mathcal{L}_{\kappa,\omega}$ differs from the finitary language $\mathcal{L}_{\omega,\omega}$ in that it allows conjunctions and disjunctions over infinite sets, with cardinality bounded by κ. Of particular importance will be $\mathcal{L}_{\omega_1,\omega}$, which allows countable conjunctions and disjunctions.

By a formula of $\mathcal{L}_{\infty,\omega}$, we mean a formula of $\mathcal{L}_{\kappa,\omega}$ for some κ. Formulas of $\mathcal{L}_{\infty,\omega}$ are closed under taking conjunctions and disjunctions of arbitrary sets. Of course, $\mathcal{L}_{\infty,\omega}$ is not a set, so in practice, we will work with fragments. This need also arises when working in an infinitary language, when one seeks to bound the cardinality of the set of formulas under consideration. For our purposes it will suffice to define a fragment $\mathbb{A}$ to be a set of formulas of $\mathcal{L}_{\infty,\omega}$ with the following properties.

1. If $\psi \in \mathbb{A}$, $\neg\psi \in \mathbb{A}$.

2. If $\psi \in \mathbb{A}$, every subformula of ψ is in $\mathbb{A}$.

We define the subformulas of ψ by recursion, as follows.

1. If ψ is atomic, ψ is the only subformula of ψ.

2. If $\psi = \neg\phi$, then ψ and the subformulas of ϕ are subformulas of ψ.

3. If $\psi = \forall\overline{y}\phi(\overline{y})$ or $\psi = \forall\overline{y}\phi(\overline{y})$, then ψ and the subformulas of ϕ are subformulas of ψ.

4. If $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$ or $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$, then ψ is a subformula of ψ, as are any subformulas of ϕ for $\phi \in \Phi$.

Starting with a formula $\psi \in \mathcal{L}_{\kappa^+,\omega}$ and closing under negations and subformulas, we obtain a fragment $\mathbb{A}$ containing ψ, of cardinality at most κ.

Formulas of $\mathcal{L}_{\kappa,\omega}$ can be coded for by labelled well founded trees, in which each node has fewer than κ extensions. We say that a formula of $\mathcal{L}_{\omega_1,\omega}$ is computable if the corresponding labelled countable tree is computable. We will define two notions of complexity for infinitary formulas. We say that an infinitary formula is Π_n if it is equivalent to a formula of the form

$$\bigwedge\!\!\!\bigwedge_{i_1} \forall\overline{x}_1 \bigvee\!\!\!\bigvee_{i_2} \exists\overline{x}_2 \ldots C_{i_n} Q\overline{x}_n \theta_{\overline{i}}(\overline{x})$$

where each $\theta_{\overline{i}}$ is a finitary quantifier free formula, C is an infinite conjunction or disjunction, and Q is a universal or existential quantifier, respectively, depending on if n is odd or even. In other words, a Π_n formula begins with an infinite conjunction of universal quantifiers, and alternates between these and infinite disjunctions of existential quantifiers n times. A Σ_n formula is one of the form

$$\bigvee\!\!\!\bigvee_{i_1} \exists\overline{x}_1 \bigwedge\!\!\!\bigwedge_{i_2} \forall\overline{x}_2 \ldots C_{i_n} Q\overline{x}_n \theta_{\overline{i}}(\overline{x})$$

We say that a formula of $\mathcal{L}_{\omega_1,\omega}$ is Σ_n^c (Π_n^c respectively) if it is computable and Σ_n (Π_n respectively). For a more detailed exposition of these classes of formulas, see [Mon21].

We will also define a less sensitive notion of complexity. We denote these classes of formulas $\forall_n$ and $\exists_n$. We define them as above, except that we do not count alternations of connectives and quantifiers together, and instead count quantifier alternations alone. As an illustration, a formula

$$\forall\overline{x} \bigvee\!\!\!\bigvee_{i} \theta_i(\overline{x})$$

is Π_2, but $\forall_1$. This notion of complexity will ultimately allow us to obtain stronger results in the conclusion of Section 3. For finitary formulas, these notions of complexity coincide, and we will not distinguish between them.

Finally, we reserve the symbols $\mathbf{\Pi}^0_n$ and $\mathbf{\Sigma}^0_n$ for the Borel hierarchy, and the lightface Π^0_n and Σ^0_n for the arithmetic hierarchy. When we describe a formula as Π^0_n or Σ^0_n, it will be a formula in the language of arithmetic, which we will call an $\mathbb{N}$-formula.

2 Relative Decidability

In this section, we will go through several results concerning relative decidability properties. We will construct examples separating a few different relative decidability properties, and characterize the property of uniform relative decidability in n Turing jumps by a syntactic condition.

2.1 Decidability of Computable Models

We will first show that the decidability of countable models of a theory does not imply relative decidability of a theory.

Theorem 2.1. *There is a c.e. theory T such that every computable model of T is decidable, but T is not relatively decidable.*

Proof. We will construct T using a tree $\mathfrak{T} \subset 2^{<\omega}$ with the following properties.

1. $\mathfrak{T}$ is computable.
2. $\mathfrak{T}$ is infinite, so has infinite paths.
3. $\mathfrak{T}$ has no computable infinite path.

For a construction of a tree with these properties, see [Bau06]. One essentially diagonalizes against every possible computable path, allowing each program to run for a bounded period of time when evaluating a finite string, corresponding to the string's length.

Let $\mathcal{L}$ be the language containing binary relation E and for each $n \in \omega$ a unary relation P_n. T has the following axioms.

1. E is a symmetric, irreflexive relation.
2. For each $\sigma \in 2^{<\omega} \setminus \mathfrak{T}$, the statement
$$(\forall x)\left(((\exists y)E(x,y)) \to \neg\left(\bigwedge_{\sigma(n)=1} P_n(x) \wedge \bigwedge_{\sigma(n)=0} \neg P_n(x)\right)\right)$$
3. For each n, the statement
$$(\forall x)((\forall y)\neg E(x,y) \to \neg P_n(x))$$

Any model $\mathcal{A}$ of T is then a graph in which every vertex is labelled by an infinite binary sequence, whose nth digit is given by P_n. If a vertex is not connected to any other vertex, its sequence is the constant sequence with value 0, and if a vertex is connected to another vertex, its sequence is a path through $\mathfrak{T}$.

For any vertex in $\mathcal{A}$, its sequence can be computed from $D(\mathcal{A})$, so if $\mathcal{A}$ is computable, these sequences must be computable. Because $\mathfrak{T}$ has no computable paths, $\mathcal{A}$ must have no edges. Any instance of a formula $\varphi(\overline{x})$ can then be decided by replacing all instances of E and P_n with the truth value false, reducing the formula to a formula in the language containing only $=$.

On the other hand, because $\mathfrak{T}$ has plenty of noncomputable paths, any graph $\mathcal{G}$ is the reduct of some model of T. Let $\pi \in 2^\omega$ be a path through $\mathfrak{T}$. We can construct a model of T which is not relatively decidable as follows. Consider the graph $\mathcal{G}$ containing for each natural number n, vertices u_n and v_n, as well a distinct vertex w. Every vertex u_n and v_n is connected to w, and all vertices are labelled with π. Consider the formula $\varphi(x) = (\exists y)(\exists z)(y \neq z \wedge E(x, y) \wedge E(x, z))$. We will construct $\mathcal{G}$ so that the set $Y = \{n | \mathcal{G} \models \varphi(u_n)\}$ is not computable from $D(\mathcal{G})$. This set is computable from $E(\mathcal{G})$, so this will imply that G is not relatively decidable.

Let Φ_e^X be the eth program relative to a set X. We add edges to $\mathcal{G}$ in stages so as to satisfy the following requirements.

R_e: $\Phi_e^{D(\mathcal{G})}(e) \neq Y(e)$

We rank these by priority in order $R_1, R_2, \ldots$ We say that R_e requires attention at stage s if $\Phi_e^{D(\mathcal{G}_s)}(n_e)$ halts in less than s steps. At a given step s, let R_e be the highest priority requirement among the first s that requires attention, if one exists. Reserve the part of the diagram used in the computation of $\Phi_e^{D(\mathcal{G}_s)}(e)$. If $\Phi_e^{D(\mathcal{G}_s)}(e)$ halts with output 0, then choose m, large enough so that $E(u_e, v_m)$ and $\neg E(u_e, v_m)$ are outside of any reserved part of the diagram, and outside of the first e sentences with constants naming the elements of $\mathcal{G}$. Add an edge connecting u_e to v_m, adding e to Y_{s+1}. If it halts with any other value, add no edges.

In the limit, all requirements will be satisfied, so $\mathcal{G}$ will fail to be relatively decidable. The $\mathcal{G}$ constructed above has diagram computable in π. A path through $\mathfrak{T}$ can be computed in $0'$, so T has models which fail to be relatively decidable that are computable in $0'$.

□

Elaborating on this, we can obtain a complete theory with this property.

Theorem 2.2. *There is a complete, computable theory T such that every computable model of T is decidable, but T is not relatively decidable.*

Proof. We will use $\mathfrak{T}$ as in the previous theorem. Let $\mathcal{L}$ be the language containing a binary relation R, unary relations U and Q, and for each n unary relations P_n and L_n. T has the following axioms:

1. $(\forall x)(\forall y)(R(x, y) \rightarrow (U(x) \wedge \neg U(y)))$
2. For each $m \geq n$, the sentence $(\forall x \in U)(L_m(x) \rightarrow L_n(x))$
3. For each $\sigma \in 2^{<\omega} \setminus \mathfrak{T}$ of length n, the sentence
$$(\forall x \in U)\left(L_n(x) \rightarrow \neg\left(\bigwedge_{\sigma(i)=1} P_i(x) \wedge \bigwedge_{\sigma(i)=0} \neg P_i(x)\right)\right)$$
4. $(\forall y \notin U)(\exists! x \in U)(R(x, y))$
5. For each n, the sentences
$$(\forall x \in U)(L_n(x) \rightarrow (\exists_{\geq n} y \notin U)R(x, y))$$
and
$$(\forall x \in U)(\neg L_n(x) \rightarrow (\exists_{<n} y \notin U)R(x, y))$$
6. $(\forall x \in U)(\exists_{<2} y \notin U)(R(x, y) \wedge Q(y))$
7. $(\forall x \in U)\neg Q(x)$
8. For each n, the sentence $(\forall y \notin U)(\neg L_n(x) \wedge \neg P_n(x))$
9. Any finite combination of L_n and P_n consistent with axioms (2) and (3) is realized by infinitely many $x \in U$. For infinitely many of these, there is some y with $R(x, y)$ and $Q(y)$, and for infinitely many, there is no such y.

Any model of T consists of an infinite set U, each element x of which is labelled with an infinite binary sequence whose nth digit is given by $P_n(x)$. Each $x \in U$ satisfies either all L_n up to L_N for some N, and none above it, in which case we say that it has length N, or satisfies all L_n, in which case we say that it has infinite length. If x has length n, the first n digits of its sequence are a string in $\mathfrak{T}$, and if it has infinite length, its sequence is a path through $\mathfrak{T}$. For each $x \in U$, there is an associated subset of U^c, whose cardinality is given by the length of x. These subsets partition U^c, and each contains at most one element satisfying Q.

We will first show that T is satisfiable. Any finite subset S of T contains only finitely many instances of axiom scheme (9), and uses finitely many of the P_n and L_n. This finite subset asserts that some finite combinations of these L_n and P_n consistent with axioms (2) and (3) are realized at some finite number of times. We can construct a finite structure realizing this, containing an element of U for instance of each combination we need to realize, of length no more than prescribed by the L_i. Each of these can be labelled by the prescribed combination of the P_i, and we interpret all unused P_i as empty. For each element of U, we have a set of elements of U^c of cardinality given by the element's length. If the combination we are realizing is prescribed to have an associated element satisfying Q, we let the first element of its associated subset satisfy Q. This is then a model of S. The compactness theorem implies that T is satisfiable.

We will now show that T is computable and complete. Because the axioms are a computable set of sentences, it suffices to show that is complete. For this, it suffices to show that for any finite subset $\mathcal{F} \subset \mathcal{L}$, any countable models $\mathcal{A}$ and $\mathcal{B}$ of T are elementarily equivalent as $\mathcal{F}$-structures. We will show that any such $\mathcal{A}$ and $\mathcal{B}$ have isomorphic elementary extensions, as $\mathcal{F}$-structures. It suffices to consider $\mathcal{F}$ of the form

$$\{U, L_0, \ldots, L_N, P_0, \ldots, P_N, Q, R\}$$

Let σ be a length N binary sequence in $\mathfrak{T}$. In both $\mathcal{A}$ and $\mathcal{B}$, there are infinitely many elements of U, of arbitary length, such that the first N digits of the binary sequence labelling them is σ. Infinitely many of these have an associated element satisfying Q, and infinitely many do not. By the compactness theorem, there are countable elementary extensions $\mathcal{A}'$ and $\mathcal{B}'$ of $\mathcal{A}$ and $\mathcal{B}$ with the following properties. Each contains infinitely many $x \in U$ for which there are infinitely many $y \notin U$ satisfying $R(x, y)$. For a binary sequence of length N in $\mathfrak{T}$, infinitely many of these xs are labelled with it. For infinitely many of these, one of the associated y satisfies Q, and for infinitely many, none of the associated y satisfy Q.

We can then construct partial $\mathcal{F}$-isomorphisms between $\mathcal{A}'$ and $\mathcal{B}'$, by mapping an $x \in U^{\mathcal{A}}$ to an x' in $U^{\mathcal{B}}$ of the same length, labelled with the same string, and which has an associated y satisfying Q if and only if x does. The ys associated to x can be mapped to y's associated to x'. We can similarly construct partial isomorphisms from $\mathcal{B}'$ to $\mathcal{A}'$. We have countably many options to chose from in each structure, so we can use the back and forth method to construct an isomorphism $\mathcal{A}' \simeq \mathcal{B}'$. We conclude that T is a computable, complete theory. The effective completeness theorem then implies that T has decidable models.

We will now show that any computable model $\mathcal{A}$ of T is decidable. The binary sequence assigned to any $x \in U$ can be computed from $D(\mathcal{A})$, so because $\mathfrak{T}$ has no computable paths, each such x must have finite length. Let $a \in \mathcal{A}$. We define the neighborhood of a, $N(a)$, as follows. If $a \in U^{\mathcal{A}}$, $N(a) = \{a\} \cup \{b | \mathcal{A} \models R(a, b)\}$. If $a \in \mathcal{A} \setminus U^{\mathcal{A}}$, there is a unique $b \in U^{\mathcal{A}}$ such that $\mathcal{A} \models R(b, a)$. We let $N(a) = N(b)$. For $\overline{a}$ a tuple in $\mathcal{A}$, we define $N(\overline{a})$ to be the union of $N(a_i)$ over each entry a_i of $\overline{a}$. Because each element of $U^{\mathcal{A}}$ has finite length, $N(\overline{a})$ is a finite subset of $\mathcal{A}$ for any tuple $\overline{a}$. We can think of $N(\overline{a})$ as a finite substructure of $\mathcal{A}$. Consider the atomic diagram $D(N(\overline{a}))$.

Any model of $D(N(\overline{a})) + T$ is a model of T containing $N(\overline{a})$ as a substructure. For any countable such models, and finite fragment $\mathcal{F}$ of the language, we can as above construct isomorphic elementary extensions of the models as $\mathcal{F}$-structures. This shows that $T(\overline{a})$, the set of consequences of $D(N(\overline{a})) + T$ is a complete theory. For any element of $\mathcal{A} \setminus U^{\mathcal{A}}$, we can effectively search for the unique element of $U^{\mathcal{A}}$ to which it is associated. We can effectively compute the length of an element of $U^{\mathcal{A}}$, and effectively search for elements of its neighborhood. As such, $N(\overline{a})$ can be computed uniformly in $\overline{a}$, so $D(N(\overline{a}))$ can be as well. Therefore, $T(\overline{a})$ is uniformly computable in $\overline{a}$, so for any formula $\varphi(\overline{x})$ and tuple $\overline{a}$, we can computably decide if $\mathcal{A} \models \varphi(\overline{a})$. Consequently, $E(\mathcal{A})$ computable, so $\mathcal{A}$ is decidable.

We will now show that T has models which are not relatively decidable. Let $\mathcal{A}$ be any countable model of T. We will construct a model $\mathcal{A}^*$ that is not relatively decidable, which extends $\mathcal{A}$. Add countably many new elements, $u_1, u_2, \dots$ to $\mathcal{A}$, which will satisfy U in $\mathcal{A}^*$. For each u_n, add countably many new elements $v_{n,1}, v_{n,2}, v_{n,3}, \dots$, which will satisfy $\neg U$. Extend R by letting $\mathcal{A}^* \models R(u_n, v_{n,i})$ for each i. Let each u_n have infinite length, and label all of them with some path π in $\mathfrak{T}$. Because $\mathcal{A} \models T$, $\mathcal{A}^*$ as constructed so far is a model of T. Let $\varphi(x)$ be the formula $(\exists y)(R(x,y) \wedge Q(y))$. Let Y be the set $\{n | \mathcal{A}^* \models \varphi(u_n)\}$. We will define the relation Q on the new elements $v_{n,i}$ so as to make Y not computable from $D(\mathcal{A}^*)$, which will imply that $\mathcal{A}^*$ is not relatively decidable, because Y is computable from $E(\mathcal{A}^*)$. Let $\mathcal{A}^*_s$ and Y_s be $\mathcal{A}^*$ and Y as defined at the sth stage of the construction. Let Φ_e^X be the eth program relative to a set X. We will satisfy the requirements

R_e: $\Phi_e^{D(\mathcal{A}^*)}(e) \neq Y(e)$

which we rank by priority $R_1, R_2, R_3, \dots$. As in the previous theorem, we say that R_e requires attention at step s if $\Phi_e^{D(\mathcal{A}^*_s)}(e)$ halts in less than s steps. At step s, let R_e be the requirement among the first s of highest priority that requires attention, if one exists. Reserve the part of the diagram used in the computation of $\Phi_e^{D(\mathcal{A}^*_s)}(e)$. If $\Phi_e^{D(\mathcal{A}^*_s)}(e)$ halts with value 0, choose i large enough so that $Q(v_{e,i})$ and $\neg Q(v_{e,i})$ lie outside of any part of $D(\mathcal{A}^*)$ that has been reserved so far, and outside of the first e sentences with constants naming the elements of $\mathcal{A}^*$. Let $\mathcal{A}^*_{s+1} \models Q(v_{e,i})$, adding e to Y_{s+1}.

In the limit, all requirements are satisfied, so Y is not computable from $D(\mathcal{A}^*)$. For each e, there is at most one i such that $\mathcal{A}^* \models Q(v_{e,i})$, so $\mathcal{A}^*$ is still a model of T, which fails to be relatively decidable. The model $\mathcal{A}^*$ constructed above is computable in $D(\mathcal{A}) \oplus \pi$. Taking $\mathcal{A}$ to be a decidable model of T, and π to be a path computable in $0'$, $\mathcal{A}^*$ is computable in $0'$.

□

A variation on the construction of the previous theory can be used to prove a lower bound on the complexity of the property of a theory that all its computable models are decidable. We will need the following result.

Theorem 2.3. *Let T_e be the eth computable subtree of $2^{<\omega}$. The set $\{e | T_e$ has a computable path$\}$ is Σ^0_3 hard.*

Proof. Let $\varphi(w) = \exists x \forall y \exists z \theta(w,x,y,z)$ be a Σ^0_3 predicate. Because $\{e | W_e$ is infinite$\}$ is Π^0_2 complete, there is a computable function f such that $W_{f(w,x)}$ is infinite if and only if $\forall y \exists z \theta(w,x,y,z)$.

For $\sigma_0\sigma_1 \dots \sigma_\ell = \sigma \in 2^{<\omega}$, let σ^* be the string $\sigma_0 1 \sigma_1 1 \dots 1 \sigma_\ell 1$ Let $\mathfrak{T}^*$ be the tree obtained by closing the set $\{\sigma^* | \sigma \in \mathfrak{T}\}$ under initial segments. $\mathfrak{T}^*$ is computable because $\mathfrak{T}$ is, and any path in $\mathfrak{T}^*$ computes a path in $\mathfrak{T}$, so $\mathfrak{T}^*$ has no computable path. We say that a string τ is a branch off of $\mathfrak{T}^*$ if τ is of the form $\sigma_0 1 \sigma_1 1 \dots 1 \sigma_\ell 0$ for $\sigma \in \mathfrak{T}$. Because $\mathfrak{T}$ is computable, we can effectively enumerate the branches off $\mathfrak{T}^*$ in lexicographic order. Let τ_n be the nth branch off $\mathfrak{T}^*$.

For e a natural number, we will construct a computable tree $T_{g(e)}$ such that $T_{g(e)}$ has a computable path if and only if $\varphi(e)$ is true. We construct $T_{g(e)}$ in stages, as follows. $T_{g(e)}$ contains $\mathfrak{T}^*$ as a subtree. For each n, we construct a copy of $\mathfrak{T}$ above τ_n, one layer at a time. If at stage s, a new element is enumerated into $W_{f(e,n)}$, we stop building the copy of $\mathfrak{T}$ over τ_n at the sth layer, and begin building a new one over the leftmost node in this layer.

Suppose $\varphi(e)$ is true. Then, for some n, infinitely many elements are enumerated into $W_{f(e,n)}$. Our construction of trees over τ_n will then be interrupted infinitely often, so $T_{g(e)}$ has exactly one path going through τ_n, which passes through the leftmost node at each stage we are interrupted. Because $T_{g(e)}$ is computable, the unique path through τ_n can be computed.

Suppose $\varphi(e)$ is false. A path through $T_{g(e)}$ is either a path through $\mathfrak{T}^*$, which is uncomputable, or passes through some branch τ_n off of $\mathfrak{T}^*$. In this case, $\forall y \exists z \theta(e,n,y,z)$ is false, so $W_{g(n,e)}$ is finite. Therefore, all paths passing through τ_n are, up to a finite initial segment, paths though a copy of $\mathfrak{T}$, so are uncomputable.

□

Theorem 2.4. *The set $\{e |$ every computable model of the eth computable complete theory is decidable$\}$ is Π^0_3 hard.*

Proof. Let φ be as in the previous theorem. We will construct a complete computable theory Th_e, uniformly in e, such that every computable model of Th_e is decidable if and only if $\varphi(e)$ is false. We construct Th_e exactly as in Theorem 2, with $T_{g(e)}$ replacing $\mathfrak{T}$. $T_{g(e)}$ shares all the properties of $\mathfrak{T}$ that imply that the resulting theory is computable and complete, namely, that it is computable and infinite.

As in Theorem 2, we see that if $\varphi(e)$ is false, them $T_{g(e)}$ has no computable path, so every computable model of Th_e is decidable. If $\varphi(e)$ is true, then $T_{g(e)}$ has a computable path. In Theorem 2, we showed that the theory has a model which is not relatively decidable that is computable relative to a path in the tree. In this case, this implies that Th_e has a computable model which is not decidable.

□

2.2 Uniform Relative Decidability in n Turing Jumps

Recall that a theory T is uniformly relatively decidable if there is a Turing functional Φ such that for any countable $\mathcal{A} \models T$, $\Phi^{D(\mathcal{A})}$ computes $E(\mathcal{A})$, the elementary diagram of $\mathcal{A}$. That is, Φ represents a uniform procedure for deciding formulas in models of T relative to their diagrams. Chubb, Miller, and Solomon gave the following characterization of uniformly relatively decidable c.e. theories.

Theorem 2.5. *[CMS21] The following are equivalent, for T a c.e. theory.*

1. *T is model complete.*
2. *T is uniformly relatively decidable*
3. *T has effective quantifier elimination down to Σ_1^c formulas.*
4. *T has effective quantifier elimination down to finitary Σ_1 formulas*

This theorem shows that uniform relative decidability always has a structural reason for showing up, in the form of a quantifier elimination property. We can extend this theorem, characterizing weaker quantifier elimination properties in terms of weaker computability properties, as follows.

Theorem 2.6. *The following are equivalent, for T a c.e. theory.*

1. *There is a Turing functional Φ such that for any countable $\mathcal{A} \models T$, $\Phi^{D(\mathcal{A})^{(n)}}$ computes $E(\mathcal{A})$.*
2. *T has effective quantifier elimination down to Σ_{n+1}^c formulas.*
3. *T has effective quantifier elimination down to finitary Σ_{n+1} formulas.*

We will need two main tools to prove this theorem. The first is the following theorem, which provides a connection between relative computability and definability in infinitary logic. We say that a relation $R \subset \mathcal{A}^n$ is uniformly relatively intrinsically Σ_n^0 if there is a Σ_n^0 $\mathbb{N}$-formula $\psi(\overline{x}, \dot{X})$ such that $\psi(\overline{x}, D(\mathcal{B}))$ defines $R^{\mathcal{B}} \subset \mathcal{B}^n (= \omega^n)$ for any $\mathcal{B} \simeq \mathcal{A}$.

Theorem 2.7 (Ash, Knight, Manasse, Slaman, and Chisholm)**.** *[AKMS89][Chi90] A relation $R \subset \mathcal{A}^n$ is uniformly relatively intrinsically Σ_n^0 if and only if R is definable by a Σ_n^c in $\mathcal{A}$. In fact, if R is defined in every copy $\mathcal{B} \simeq \mathcal{A}$ by a Σ_n^0 $\mathbb{N}$-formula $\psi(\overline{y}, D(\mathcal{B}))$, then R is definable in $\mathcal{A}$ by a Σ_n^c formula*

$$\lambda(\overline{x}) = \exists \overline{q} \bigvee_{\overline{k} \in |\overline{q}|^n} \left(\overline{x} = \overline{q}(\overline{k}) \wedge \mathrm{Force}_{\psi(\overline{k}, \dot{X})}(\overline{q}) \right)$$

where $\exists \overline{q}$ is shorthand for an infinite disjunction over all possible lengths of $\overline{q}$, and $\mathrm{Force}_{\psi(\overline{k}, \dot{X})}$ *is a Σ_n^c formula uniformly computable in ψ.*

For the details of the formula $\lambda(\overline{x})$, see [Mon21], Chapter 7. For our purposes, all we need to know about it is that it does not depend on the structure $\mathcal{A}$ and that it is uniformly computable in ψ. The second tool we will need bridges the gap between definability in infinitary logic and definability in finitary logic, while preserving the complexity of formulas.

Proposition 2.8. *Let T be a finitary theory. Let ψ be an infinitary Π_n formula which is equivalent to a finitary formula φ in all models of T. Then, ψ and φ are equivalent to a finitary Π_n formula in all models of T.*

In Section 3, we will prove a number of special cases of this proposition, culminating in a proof of a more general statement. With these tools, we can proceed with our proof of Theorem 2.6.

Proof. We will first show that (1) implies (2). If T satisfies (1), and $\varphi(\overline{x})$ is a (finitary) formula in the language of T in m free variables, φ defines a relation $R \subset \mathcal{A}^m$ for any countable $\mathcal{A} \models T$. R is uniformly computable from $E(\mathcal{A})$, so is uniformly computable from $D(\mathcal{A})^{(n)}$. By Post's theorem, R is thus uniformly relatively intrinsically Σ^0_{n+1}, so by Theorem 2.7, is definable in $\mathcal{A}$ by the Σ^c_{n+1} formula

$$\lambda(\overline{x}) = \exists \overline{q} \bigvee_{\overline{k} \in |\overline{q}|^n} \left(\overline{x} = \overline{q}(\overline{k}) \wedge \mathrm{Force}_{\psi(\overline{k}, \dot{X})}(\overline{q})\right)$$

We then have that λ is equivalent to φ in any model $\mathcal{A} \models T$. Moreover, the $\mathbb{N}$-formula ψ defining R can be computed uniformly from φ, and the Turing functional Φ. As such, T has effective quantifier elimination to infinitary Σ^c_{n+1} formulas.

We will now show that (2) implies (3). If T has quantifier elimination to infinitary Σ^c_{n+1} formulas, applying Proposition 2.8 to each formula (using negation to switch between Σ_{n+1} and Π_{n+1}) implies that T has quantifier elimination to finitary Σ_{n+1} formulas. Moreover, because T is c.e. we can for any φ effectively search for a Σ_{n+1} formula and a proof of its equivalence to φ from T, so this quantifier elimination is effective.

Finally, we will show that (3) implies (1). In order to compute $E(\mathcal{A})$, we need to decide formulas in $\mathcal{A}$. To decide a formula φ, we first find Σ_{n+1} formulas $\exists \overline{y} \theta(\overline{xy})$ and $\exists \overline{y} \eta(\overline{xy})$ equivalent to φ and $\neg\varphi$ respectively over T, where θ and η are Π_n. To decide if $\mathcal{A} \models \varphi(\overline{a})$, we search tuples $\overline{b}$ and $\overline{c}$ of elements of $\mathcal{A}$, checking if $\mathcal{A} \models \theta(\overline{ab})$ or if $\mathcal{A} \models \eta(\overline{ac})$ using $D(\mathcal{A})^{(n)}$. One of the Σ_{n+1} formulas is true, so this search will eventually terminate, at which point we can determine if $\mathcal{A} \models \varphi(\overline{a})$. $\square$

2.3 Computability of Fragments of the Diagram

Before embarking towards a proof of Proposition 2.8, we will demonstrate a few more distinctions between relative decidability properties.

For $\mathcal{A}$ a structure, $E_{\Pi_n}(\mathcal{A})$ ($E_{\Sigma_n}(\mathcal{A})$) is the set of Π_n (Σ_n) sentences in the elementary diagram of $\mathcal{A}$. For any structure, $E_{\Sigma_n}(\mathcal{A}) \equiv_T E_{\Pi_n}(\mathcal{A})$ uniformly, so (uniform) computability of one is equivalent to (uniform) computability of the other. We will switch between them as convenient. We will use the convention that $0^{(n)}$ is represented by $E_{\Sigma_n}(\mathbb{N})$.

Theorem 2.9. *Let T be a c.e. theory. Suppose there is a Turing functional Φ such that for $\mathcal{A}$ a countable model of T, $\Phi^{D(\mathcal{A})}$ computes $E_{\Pi_1}(\mathcal{A})$. Then T is uniformly relatively decidable.*

Proof. It suffices to show that T has quantifier elimination to Σ_1 formulas. We will prove by induction of n that any Π_n formula is equivalent over T to a Σ_1 formula.

Let $\varphi(\overline{x})$ be a Π_1 formula with n free variables. For $\mathcal{A}$ a countable model of T, the set $R \subset \mathcal{A}^n$ defined by φ can be uniformly computed from $E_{\Pi_1}(\mathcal{A})$, and so can be uniformly computed from $D(\mathcal{A})$. Therefore, R is uniformly relatively intrinsically c.e, so by the Ash-Knight-Manasse-Slaman-Chisholm theorem, is defined in $\mathcal{A}$ by a Σ^c_1 formula $\psi(\overline{x})$, which is independent of the structure $\mathcal{A}$. That is, φ and ψ are equivalent in any countable model of T, so by the downward Löwenheim-Skolem theorem, in any model of T.

Suppose $\mathcal{A}$ and $\mathcal{B}$ are models of T, $\mathcal{A} \subset \mathcal{B}$, and $\overline{a} \subset \mathcal{A}$ such that $\mathcal{A} \models \varphi(\overline{a})$. Then, $\mathcal{A} \models \psi(\overline{a})$, so because ψ is Σ^c_1, $\mathcal{B} \models \psi(\overline{a})$, which implies that $\mathcal{B} \models \varphi(\overline{a})$. Because this holds for any models $\mathcal{A}$ and $\mathcal{B}$ of T, φ is equivalent to a finitary Σ_1 formula η over T.

Suppose that any Π_n formula is equivalent over T to a Σ_1 formula. Let φ be a Π_{n+1} formula. Then, $\neg\varphi(\overline{x})$ is equivalent to a Σ_{n+1} formula $\exists \overline{y} \theta(\overline{x}, \overline{y})$, for θ a Π_n formula. Let η be a Σ_1 equivalent to θ. Then, $\neg\varphi(\overline{x})$ is equivalent to $\exists \overline{y} \eta(\overline{x}, \overline{y})$, which is Σ_1. Hence, φ is equivalent to a Π_1 formula, which is in turn equivalent to a Σ_1 formula. We conclude that every formula is equivalent over T to a Σ_1 formula, so T is uniformly relatively decidable.

$\square$

The analogous statement for copies of a single structure is not true in general.

Theorem 2.10. *There is a countable structure $\mathcal{A}$ and a Turing functional Φ such that $\Phi^{D(\mathcal{B})}$ computes $E_{\Sigma_1}(\mathcal{B})$ for any $\mathcal{B} \simeq \mathcal{A}$, but $\mathcal{A}$ is not relatively decidable.*

Proof. Let $\mathcal{L}$ consist of the language of arithmetic $\{0, 1, +, \times, <\}$, together with new constants c_n for each $n \in \omega$, and new unary relations P, W_0, and W_1. Let the domain of $\mathcal{A}$ be the disjoint union of two copies of ω, ω_0, the interpretation of W_0, and ω_1, the interpretation of W_1. On ω_0, let the symbols in the language of arithmetic have the standard interpretation. Let the c_n label the elements of ω_1, and let P be interpreted as $\{c_n | n \in 0'\}$. If $a \in \omega_1$, then for any b, let $a + b = a \times b = b + a = b \times a = c_0$.

Consider a fixed $\mathcal{B} \simeq \mathcal{A}$. Any Σ_1 formula in m free variables $\varphi(\overline{x}) = \exists y_1, \ldots, y_n \theta(\overline{x}, \overline{y})$, for θ quantifier free, is equivalent to the formula

$$\bigvee_{\sigma \in 2^m} \left(x_1 \in W_{\sigma(1)} \wedge \cdots \wedge x_m \in W_{\sigma(m)} \wedge \bigvee_{\sigma' \in 2^n} \exists y_1 \in W_{\sigma'(1)} \ldots \exists y_n \in W_{\sigma(n)} \theta(\overline{x}, \overline{y}) \right)$$

so to decide an instance of $\varphi(\overline{x})$, it suffices to decide each existential disjunct of the above formula. Writing θ in disjunctive normal form, it suffices to consider the case in which θ is a conjunction of literals. For each term $\tau(\overline{x}, \overline{y})$, whether τ is in ω_0 or ω_1 depends only on its syntactic form and on whether the x_i and the y_i lie in ω_0 and ω_1. A term in ω_1 is either a variable ranging over ω_1, a constant c_i, or an expression always equal to c_0, which can be replaced with c_0. We may thus assume such terms contain only variables ranging over ω_1. Likewise, terms in ω_0 contain only variables ranging over ω_0.

Atomic sentences of the form $W_1(\tau)$ and $W_2(\tau)$ can be replaced with their truth values. An atomic sentence $P(\tau)$ can be replaced with the truth value false whenever $\tau \in \omega_0$, as can atomic sentences involving $<$ with terms in ω_1 and equations between terms in different parts of the structure. As such, it suffices to consider θ of the form

$$\rho(\overline{x}, \overline{y}) \wedge \lambda(\overline{x}, \overline{y})$$

where ρ is in the language of arithmetic and contains only variables ranging over ω_0, and λ is in the language $\{P, c_1, c_2, \ldots\}$ and contains only variables ranging over ω_1. Because they contain disjoint sets of variables, the quantified formula is equivalent to the conjunction of the quantifiers applied to ρ and to λ separately, so it suffices to decide each of these. The case of λ can be decided because the theory of W_1 in the language language $\{P, c_1, c_2, \ldots\}$ is complete, as any two countable models have isomorphic elementary extensions, each containing countably many elements not equal to any c_i satisfying P, and countably many satisfying $\neg P$. This theory is axiomatized by the atomic sentences in $D(\mathcal{B})$, so is computable from it.

In the case of ρ, we need to decide a Σ^0_1 arithmetic formula, plugging in elements $\overline{b}$ in $\mathcal{B}$ for $\overline{x}$. We may assume all of $\overline{b}$ satisfy W_0, so lie in the copy of ω_0. We can effectively search for terms $1 + \cdots + 1$ equal to each b_i, and thus turn the formula into a Σ^0_1 sentence in the language of arithmetic. We can decide this sentence by computing its Gödel number i, and checking if $P(c_i) \in D(\mathcal{B})$. This describes a uniform procedure to compute $E_{\Sigma_1}(\mathcal{B})$ from $D(\mathcal{B})$.

$\mathcal{A}$ is not relatively decidable, because

$$E(\mathcal{A}) \geq_T \mathrm{Th}(\mathbb{N}) >_T 0' \geq_T D(\mathcal{A})$$

$\square$

We can replace $0'$ with $0^{(n)}$ in the above construction. In order to decide a Σ_n formula, we carry out the same procedure of converting it to prenex normal form, then considering all possible combinations of W_i that the variables can lie in. This reduces the problem of deciding the formula, as above, to deciding a Σ^0_n sentence about W_0 in the language of arithmetic, for which we use $0^{(n)}$, and deciding a sentence in the language $\{P, c_1, c_2 \ldots\}$ about W_1. Using the fact that $\mathrm{Th}(\mathbb{N}) >_T 0^{(n)}$ for any n, we have the following generalization.

Theorem 2.11. *For any n, there is a countable structure $\mathcal{A}$ and Turing functional Φ such that $\Phi^{D(\mathcal{B})}$ computes $E_{\Sigma_n}(\mathcal{B})$ for any $\mathcal{B} \simeq \mathcal{A}$, but $\mathcal{A}$ is not relatively decidable.*

Using a variation on the above construction, we can realize all of these properties in the same structure.

Theorem 2.12. *There is a countable structure $\mathcal{A}$ such that for each n, there is a Turing functional Φ_n such that $\Phi_n^{D(\mathcal{B})}$ computes $E_{\Sigma_n}(\mathcal{B})$ for $\mathcal{B} \simeq \mathcal{A}$, but $\mathcal{A}$ is not relatively decidable.*

Proof. The idea is as follows. We will include a copy of $0^{(n)}$ in the structure for each n, each of which can be read off of the diagram using a distinct "key." We let $\mathcal{L}$ consist of the language of arithmetic, together with new constants $c_{i,k}$ for natural numbers i, k, and unary relations W_0, W_1 and P. $\mathcal{A}$ consists of the disjoint union of two sets. One is a copy of ω, the interpretation of W_0, on which the language of arithmetic has its standard interpretation. The other is a copy of ω^2, the interpretation of W_1.

The elements of this are labelled by the constants $c_{i,k}$. For each n, we will define a distinct k_n. We will also define a set $N \subset W_1$, which can thought of as "noise" that we add to the diagram, disjoint from the rows corresponding to the k_n. We interpret P as $\{c_{i,k_n} | i \in 0^{(n)}\} \cup N$. We extend the language of arithmetic to W_1 as in the previous theorem, so that if $a \in W_1$, $a + b = b + c = a \times b = b \times a = c_{0,0}$ for any b.

For any n, we can decide Σ_n formulas in any $\mathcal{B} \simeq \mathcal{A}$ exactly as in the previous theorem, with one modification. When we compute the Gödel number i of a Σ^0_n arithmetic sentence we need to decide, we check if $P(c_{i,k_n}) \in D(B)$. Using k_n, we can thus define a uniform procedure to compute $E_{\Sigma_n}(\mathcal{B})$ from $D(\mathcal{B})$ for any $\mathcal{B} \simeq \mathcal{A}$.

We will now show that N and the k_n can be chosen so that $\mathcal{A}$ is not relatively decidable. In particular, we will show that they can be chosen so that $\mathrm{Th}(\mathbb{N})$ is not computable from $D(\mathcal{A})$. Let Φ_e^X be the eth Turing functional relative to a set. We will define the k_n and N in stages. Let $D(\mathcal{A}_s)$ be the diagram of $\mathcal{A}$ at stage s. At step 0, set $k_1 = 1$, N to be the empty set. At step s, we will have defined $k_1 < k_2 < \cdots < k_s$, and put only finitely many elements into N. $D(\mathcal{A}_s)$ is thus computable from $0^{(s)}$, so $\mathrm{Th}(\mathbb{N})$ is not computable from it. We will show that there must be an i such that one of the following holds.

1. $\Phi_s^{D(\mathcal{A}^*)}(i) \uparrow$ for any $\mathcal{A}^*$ obtained from $\mathcal{A}_s$ by adding elements to P in rows above k_s.

2. There is an i, and a finite extension of P above row k_s so that $\Phi_s^{D(\mathcal{A}^*)}(i) \downarrow$ and $\Phi_s^{D(\mathcal{A}^*)}(i) \neq \mathrm{Th}(\mathbb{N})(i)$.

If for all i, both (1) and (2) fail, then we can compute $\mathrm{Th}(\mathbb{N})$ from $D(\mathcal{A}_s)$ as follows. For any i, we search finite extensions P above k_s until we find one so that $\Phi_s^{D(\mathcal{A}^*)}(i)$ converges, which must exist because (1) fails. The computation must have output $\mathrm{Th}(\mathbb{N})(i)$, because (2) fails. This is a contradiction, so there must be some i satisfying (1) or (2).

If i satisfies (1), then because in the remainder of the construction we will only add elements to P above row k_s, this ensures that $\Phi_s^{D(\mathcal{A})}$ will not compute $\mathrm{Th}(\mathbb{N})$. We can then choose k_{s+1} to be $k_s + 1$, and move to step $s + 1$. If i satisfies (2), then we extend N to realize the finite extension of P so that $\Phi_s^D(\mathcal{A}^*)(i) \neq \mathrm{Th}(\mathbb{N})(i)$, and choose k_{s+1} to be large enough so that changes to rows $\geq k_{s+1}$ do not change the part of the diagram used in the computation of $\Phi_s^{D(\mathcal{A}^*)}(i)$. This ensures that $\Phi_s^{D(\mathcal{A})}$ will not compute $\mathrm{Th}(\mathbb{N})$.

In the limit, we have defined $\mathcal{A}$ so that for all e, $\Phi_e^{D(\mathcal{A})}$ does not compute $\mathrm{Th}(\mathbb{N})$. On the other hand, $\mathrm{Th}(\mathbb{N})$ is computable from $E(\mathcal{A})$, so this implies that $\mathcal{A}$ is not relatively decidable.

□

3 Infinitary Logic

In this section, we go about proving Proposition 2.8. We will examine each case up to $n = 5$ separately, at which point we will have the necessary ideas to prove the general case.

3.1 Semantic Tests: $n = 0$ and $n = 1$

We can extract, from the definition of model completeness, the following semantic test for Π_1 formulas.

Theorem 3.1. *Let φ be a finitary formula and T be a finitary theory. The following are equivalent.*

1. *φ is equivalent to a Π_1 formula over T.*
2. *If $\mathcal{A} \subset \mathcal{B}$ are models of T, $\overline{a} \in \mathcal{A}$, and $\mathcal{B} \models \varphi(\overline{a})$, then $\mathcal{A} \models \varphi(\overline{a})$. That is, φ is preserved by substructures.*

Proof. If φ is equivalent to $\psi(\overline{x}) = \forall\overline{y}\theta(\overline{xy})$, then if $\overline{a} \in \mathcal{A}$, and $\mathcal{B} \models \varphi(\overline{a})$, we have that $\mathcal{B} \models \psi(\overline{a})$. For any $\overline{b} \in \mathcal{A}$, $\overline{b} \in \mathcal{B}$, so $\mathcal{B} \models \theta(\overline{ab})$. Then, $\mathcal{A} \models \theta(\overline{ab})$ because $\mathcal{A} \subset \mathcal{B}$. Thus, $\mathcal{A} \models \psi(\overline{a})$, so $\mathcal{A} \models \varphi(\overline{a})$.

Suppose, conversely, that φ is not equivalent to any Π_1 formula over T. Adding constants $\overline{c}$ for the free variables of φ, let Γ be the set of Π_1 consequences of $T + \varphi(\overline{c})$. If $T + \Gamma \vdash \varphi(\overline{c})$, then for some finite subset $S \subset \Gamma$, $T + S \vdash \varphi(\overline{c})$. Letting $\psi(\overline{c})$ be the conjunction of S, we would have that $\psi(\overline{x})$ is Π_1, and because $\overline{c}$ does not appear in T, $T \vdash \forall\overline{x}(\varphi(\overline{x}) \leftrightarrow \psi(\overline{x}))$, contradicting our assumption. As such, $T + \Gamma \nvdash \varphi(\overline{c})$, so there is some $\mathcal{A} \models T + \Gamma + \neg\varphi(\overline{c})$. We claim that $T + \varphi(\overline{c}) + D(\overline{A})$ is satisfiable. For any finite $R \subset D(\mathcal{A})$, let $\theta(\overline{ac})$ be the conjunction of R. If $T + \varphi(\overline{c}) + R$ is not satisfiable, $T + \varphi(\overline{c}) \vdash \forall\overline{y}\neg\theta(\overline{yc})$. Then, $\forall\overline{y}\neg\theta(\overline{yc}) \in \Gamma$, which is contradicted by the fact that $\mathcal{A} \models \Gamma$ and $\theta(\overline{ac}) \in D(\mathcal{A})$. Let $\mathcal{B} \models D(\mathcal{A}) + T + \varphi(\overline{c})$. Then, $\mathcal{A} \subset \mathcal{B}$ are models of T, $\mathcal{B} \models \varphi(\overline{c})$, and $\mathcal{A} \models \neg\varphi(\overline{c})$, so φ does not satisfy (2). $\square$

We can use this test to prove the $n = 1$ case.

Corollary 3.2. *If a finitary formula φ is equivalent to an infinitary Π_1 formula ψ in all models of T, a finitary theory, φ is equivalent to a finitary Π_1 formula over T.*

Proof. It suffices to show that ψ satisfies the criterion (2) of the previous theorem, as this implies that φ does as well. If $\psi\overline{y} = \bigwedge\!\!\!\bigwedge_i \forall\overline{y}\theta_i(\overline{xy})$, $\mathcal{A} \subset \mathcal{B}$, and $\mathcal{B} \models \psi(\overline{a})$ for $\overline{a} \in \mathcal{A}$, then for any i and $\overline{b} \in \mathcal{A}$, $\overline{b} \in \mathcal{B}$, so $\mathcal{B} \models \theta_i(\overline{ab})$. As such, $\mathcal{A} \models \theta_i(\overline{ab})$. We conclude that $\mathcal{A} \models \psi(\overline{a})$. $\square$

As we have define the classes Π_n, Π_0 formulas are already finitary, so the proposition is trivial there. However, there is a semantic test for quantifier free formulas which can be used to prove an analogous statement for quantifier free infinitary formulas. We can also observe that we could include more infinitary connectives in Π_1 formulas in the preceding corollary. We will ultimately more general results that embrace both of these observations, but for the rest of the special cases, we will retain the classes Π_n for simplicity.

Theorem 3.3. *Let φ be a finitary formula, and T be a finitary theory. The following are equivalent.*

1. *φ is equivalent to a finitary quantifier free formula over T.*
2. *$\mathcal{B}$, and $\mathcal{C}$ are models of T, $\mathcal{A} \subset \mathcal{B}$, and $\mathcal{A} \subset \mathcal{C}$, then for $\overline{a} \in \mathcal{A}$, $\mathcal{B} \models \varphi(\overline{a})$ if and only if $\mathcal{C} \models \varphi(\overline{a})$.*

Proof. We will use finitary, quantifier free sentences $\top$ and $\bot$ standing for true and false respectively in the case that φ is a sentence, for otherwise the language may lack quantifier free sentences. This is a technicality that does not add to the language's expressive power, and if the language has a constant c, we can replace these with $c = c$ and $c \neq c$ respectively. Otherwise, if φ is not equivalent to $\top$ or $\bot$ over T, there are models $\mathcal{B} \models T + \varphi$ and $\mathcal{B} \models T + \neg\varphi$, and we can take $\mathcal{A}$ to be empty.

If θ is quantifier free, $\mathcal{A} \subset \mathcal{B}$, $\mathcal{A} \subset \mathcal{C}$, $\mathcal{B}$ and $\mathcal{C}$ are models of T and $\overline{a} \in \mathcal{A}$, then $\mathcal{B} \models \theta(\overline{a})$ if and only $\mathcal{A} \models \theta(\overline{a})$ if and only if $\mathcal{C} \models \theta(\overline{a})$. If φ is equivalent over T to θ, then this implies that $\mathcal{B} \models \varphi(\overline{a})$ if and only if $\mathcal{C} \models \varphi(\overline{a})$.

Suppose conversely that φ is not equivalent to any quantifier free sentence formula T. Let Γ be the set of quantifier free consequences of $T + \varphi(\overline{c})$ and Γ' be the set of quantifier free consequences of $T + \neg\varphi(\overline{c})$, for $\overline{c}$ new constants. We claim that $\Gamma + \Gamma'$ is satisfiable. If not, then for some finite $S \subset \Gamma'$, $\Gamma + S$ is not satisfiable. Let $\theta(\overline{c})$ be the conjunction of S. Then, in models of T, $\varphi(\overline{c}) \vdash \theta(\overline{c}) \vdash \neg\neg\varphi(\overline{c})$. Because $\overline{c}$ does not appear in T, this implies that φ and θ are equivalent, contradicting our assumption. Let $\mathcal{A}'$ be a model of $\Gamma + \Gamma'$, and $\mathcal{A} \subset \mathcal{A}'$ be the substructure generated by the denotations of closed terms. Then, $D(\mathcal{A})$ is just the set of atomic sentences true in $\mathcal{A}$. We claim that $D(\mathcal{A}) + T + \varphi(\bar{c})$ and $D(\mathcal{A}) + T + \neg\varphi(\bar{c})$ are satisfiable. If not, say $D(\mathcal{A}) + T + \varphi(\bar{c})$, then for some finite subset $R \subset D(\mathcal{A})$, $R + T + \varphi(\bar{c})$ is not satisfiable. Letting η be the conjunction of R, we then have that $T + \varphi(\bar{c}) \vdash \eta$, so $\eta \in \Gamma$. This contradicts the fact that $\mathcal{A}' \models \Gamma$. The case of $\neg\varphi(\bar{c})$ follows for the same reason. We thus have $\mathcal{B} \models D(\mathcal{A}) + T + \varphi(\bar{c})$ and $\mathcal{C} \models D(\mathcal{A}) + T + \neg\varphi(\bar{c})$. This shows that φ does not satisfy (2). □

Corollary 3.4. *If ψ is infinitary, and quantifier free, φ is a finitary formula, and T is a finitary theory, such that ψ and φ are equivalent in all models of T, then φ and ψ are equivalent to a finitary quantifier free formula over T.*

Proof. It suffices to show that ψ satisfies condition (1) of Theorem 3.3, as this will imply that φ satisfies it as well. We will prove this by induction on the complexity of ψ.

If ψ is atomic, then because $\mathcal{A} \subset \mathcal{B}$ and $\mathcal{A} \subset \mathcal{C}$, $\mathcal{B} \models \psi(\bar{a})$ if and only if $\mathcal{A} \models \psi(\bar{a})$ if and only if $\mathcal{C} \models \psi(\bar{a})$. Suppose $\psi = \neg\phi$. Then, $\mathcal{B} \models \psi(\bar{a})$ if and only $\mathcal{B} \not\models \phi(\bar{a})$. Appealing to induction, this is true if and only if $\mathcal{C} \not\models \phi(\bar{a})$, or equivalently, $\mathcal{C} \models \psi(\bar{a})$.

Suppose $\psi = \bigvee\!\!\!\bigvee_{\phi \in \Phi} \phi$. Then, $\mathcal{B} \models \psi(\bar{a})$ if and only if for some $\phi \in \Phi$, $\mathcal{B} \models \phi(\bar{a})$. Appealing to induction, this is true if and only if $\mathcal{C} \models \phi(\bar{a})$ for some $\phi \in \Phi$, which is true if and only if $\mathcal{C} \models \psi(\bar{a})$. Likewise, if $\psi - \bigwedge\!\!\!\bigwedge_{\phi \in \Phi} \phi$, $\mathcal{B} \models \psi(\bar{a})$ if and only if for every $\phi \in \Phi$, $\mathcal{B} \models \phi(\bar{a})$. Appealing to induction, this is true if and only if for every $\phi \in \Phi$, $\mathcal{C} \models \phi(\bar{a})$, which is true if and only if $\mathcal{C} \models \psi(\bar{a})$.

□

3.2 Preservation by Chains: $n = 2$

The use of semantic tests to characterize complexity classes of formulas will prove to be fruitful. Originating in [Cha59] and [ŁS55], the following test characterizes Π_2 formulas.

Theorem 3.5. *Let $T \subset \mathcal{L}_{\omega,\omega}$ be a theory, and $\varphi \in \mathcal{L}_{\omega,\omega}$ be a formula. The following are equivalent.*

1. *φ is equivalent to a (finitary) Π_2 formula over T.*
2. *If $\mathcal{A}_0 \subset \mathcal{A}_1 \subset \ldots$ is a chain of models of T, and $\mathcal{A}_\omega = \bigcup_{n<\omega} \mathcal{A}_n$ is a model of T, then if $\bar{a} \in \mathcal{A}_0$, and $\mathcal{A}_n \models \varphi(\bar{a})$ for every n, we have that $\mathcal{A}_\omega \models \varphi(\bar{a})$.*

Proof. We will first show that finitary Π_2 formulas satisfy (2). Suppose $\varphi(\bar{x}) = \forall\bar{y}\exists\bar{z}\theta(\overline{xyz})$, and $\mathcal{A}_0 \subset \mathcal{A}_1 \subset \ldots$ is a chain of structures with $\mathcal{A}_\omega = \bigcup_n \mathcal{A}_n$, and $\mathcal{A}_n \models \varphi(\bar{a})$ for $\bar{a} \in \mathcal{A}_0$, for all n. Let $\bar{b} \in \mathcal{A}_\omega$. Then, $\bar{b} \in \mathcal{A}_n$ for some n. As such, there is a $\bar{c} \in \mathcal{A}_n$ such that $\mathcal{A}_n \models \theta(\overline{abc})$. Then, $\mathcal{A}_\omega \models \theta(\overline{abc})$. We conclude that $\mathcal{A}_\omega \models \varphi(\bar{a})$.

Next we will show that this characterizes Π_2 formulas. Let $\bar{c}$ be a tuple of new constants. Let Γ be the set of Π_2 sentences in the expanded language which are entailed by $T + \varphi(\bar{c})$. We will show that $T + \Gamma \vdash \varphi(\bar{c})$. Suppose $M \models T + \Gamma$, and let $\mathrm{Th}_{\Sigma_2}(M)$ be the set of Σ_2 sentences true in M. We claim that $T + \mathrm{Th}_{\Sigma_2}(M) + \varphi(\bar{c})$ is satisfiable. If it is not, then for some finite subset of $S \subset \mathrm{Th}_{\Sigma_2}(M)$, $T + S + \varphi(\bar{c})$ is not satisfiable. Let ψ be a Σ_2 equivalent of the conjunction of the elements of S. Then, $T + \varphi(\bar{c}) \vdash \neg\psi$, so $\neg\psi$ is equivalent to a Π_2 consequence of $T + \varphi(\bar{c})$, which then is an element of Γ. However, $M \models \psi$, which contradicts the assumption that $M \models \Gamma$.

Let $N \models T + \mathrm{Th}_{\Sigma_2}(M) + \varphi(\bar{c})$. Let $E_{\Pi_1}(M)$ be the set of Π_1 sentences in the elementary diagram of M. We now claim that $\mathrm{Th}(N) + E_{\Pi_1}(M)$ is satisfiable. It suffices to show that any finite set of Π_1 formulas true of elements of M can be satisfied in N. Taking the conjunction of this set, it suffices to show that any Π_1 formula $\beta(\bar{y})$ true of elements of M can be satisfied in N. In this case, $M \models \exists\bar{y}\beta(\bar{y})$, which is Σ_2 so because $N \models \mathrm{Th}_{\Sigma_2}(M)$, $N \models \exists\bar{y}\beta(\bar{y})$.

Let $N_1 \models \mathrm{Th}(N) + E_{\Pi_1}(M)$. We may take $N_1 \supset M$, interpreting the constants in $E_{\Pi_1}(M)$ as the respective elements of M. Further, $N_1 \equiv N$. We will now show that there is a structure M_1 such that $M_1 \succ M$ and $M_1 \supset N_1$. M_1 will be a model of $E(M) + D(N_1)$, where the constants used to name elements of M in $E(M)$ are identified with those used to name elements of M in $D(N_1)$. It suffices to show that for any finite subset $S \subset D(N_1)$, $E(M) + S$ is satisfiable. Taking the conjunction of S, it suffices to show that $E(M) + \alpha(\overline{d}, \overline{e})$ is satisfiable, where α is quantifier free formula, $\overline{d}$ are constants naming elements of M and so appear in $E(M)$, and $\overline{e}$ are constants naming elements of $N_1 \setminus M$. If this is not satisfiable in M, interpreting $\overline{d}$ as the respective elements $\overline{m}$ of M, then $M \models \forall\overline{y}\neg\alpha(\overline{m}, \overline{y})$, so $\forall\overline{y}\neg\alpha(\overline{d}, \overline{y}) \in E_{\Pi_1}(M)$. Because $N_1 \models E_{\Pi_1}(M)$, $N_1 \models \forall\overline{y}\neg\alpha(\overline{d}, \overline{y})$, which contradicts the fact that $\alpha(\overline{d}, \overline{e})$ occurs in $D(N_1)$. We conclude that $E(M) + \alpha(\overline{d}, \overline{e})$ is satisfiable, so $E(M) + D(N_1)$ is satisfiable. Taking a model of this, we have M_1.

Because $M_1 \succ M$ and $N_1 \equiv N$, $N_1 \models \mathrm{Th}_{\Sigma_2}(M) = \mathrm{Th}_{\Sigma_2}(M_1)$. Using this, we can iterate the constructions of M_1 and N_1 above to obtain a chain of structures

$$M \subset N_1 \subset M_1 \subset N_2 \subset M_2 \subset \dots$$

such that $N_1 \equiv N_2 \equiv N_3 \dots$ and $M \prec M_1 \prec M_2 \dots$. Consider $M_* = \cup_n M_n = \cup_n N_n$. By the elementary chain theorem, $M_* \succ M$, so $M_* \models T$. Because each $N_n \models \varphi(\overline{c})$, $M_* \models \varphi(\overline{c})$, so $M \models \varphi(\overline{c})$. We conclude that $T + \Gamma \vdash \varphi(\overline{c})$.

By the compactness theorem, $T + S \vdash \varphi(\overline{c})$, for S a finite subset of Γ. Taking a Π_2 sentence $\eta(\overline{c})$ equivalent to the conjunction of the sentences in S, we have that $T \vdash \varphi(\overline{c}) \leftrightarrow \eta(\overline{c})$. The constants $\overline{c}$ do not occur in T, so $T \vdash \forall\overline{x}(\varphi(\overline{x}) \leftrightarrow \eta(\overline{x}))$.

□

Corollary 3.6. *If a finitary formula φ is equivalent to an infinitary Π_2 formula ψ in all models of T, a finitary theory, φ is equivalent to a finitary Π_2 formula over T.*

Proof. We will show that infinitary Π_2 formulas satisfy the semantic test of the previous theorem. Let

$$\psi(\overline{x}) = \bigwedge_i \forall\overline{y} \bigvee_j \exists\overline{z}\theta_{i,j}(\overline{xyz})$$

. Suppose that $\mathcal{A}_0 \subset \mathcal{A}_1 \subset \dots$ is a chain of structures, $\overline{a} \in \mathcal{A}_0$, and that for each n, $\mathcal{A}_n \models \psi(\overline{a})$. Let $\mathcal{A}_\omega = \bigcup_n \mathcal{A}_n$. Then, for any i, and $\overline{b} \in \mathcal{A}_\omega$, $\overline{b} \in \mathcal{A}_n$ for some n. Then, there is a j and a $\overline{c}$ such that $\mathcal{A}_n \models \theta_{i,j}(\overline{abc})$. Consequently, $\mathcal{A}_\omega \models \theta_{i,j}(\overline{abc})$. We conclude that $\mathcal{A}_\omega \models \psi(\overline{a})$. □

3.3 Digression: Longer Chains

One might hope to generalize the semantic test used for the $n = 2$ case by considering other kinds of chains of structures. This approach, however, fails for the following reasons. First, if an ordered set I has cofinality ω, one can modify any ω-chain $(\mathcal{A}_n)_{n<\omega}$ to obtain an I indexed chain, by choosing a cofinal subset of I and filling in the gaps above $\mathcal{A}_n$ with copies of $\mathcal{A}_n$. This shows that any formula preserved by I indexed chains is also preserved by ω indexed chains, and so we don't get a characterization of a larger class of formulas. If we consider chains with cofinality greater than ω, they preserve too many formulas to give us a useful characterization. In particular, suitable ω_1 indexed chains preserve all $\mathcal{L}_{\omega_1,\omega}$ formulas.

A family of structures $(\mathcal{A}_\alpha)_{\alpha<\omega_1}$ is an ω_1-chain if it has the following properties.

1. For each $\alpha < \omega_1$, $\mathcal{A}_\alpha$ is countable.
2. For $\alpha \leq \beta < \omega_1$, $\mathcal{A}_\alpha \subset A_\beta$.
3. For limit $\beta < \omega_1$, $\mathcal{A}_\beta = \bigcup\limits_{\alpha<\beta} A_\alpha$.

For an ω_1-chain $(\mathcal{A}_\alpha)_{\alpha<\omega_1}$, we define $\mathcal{A}_{\omega_1}$ as $\bigcup\limits_{\alpha<\omega_1} A_\alpha$. $\mathcal{A}_{\omega_1}$ may be uncountable.

Theorem 3.7. *Suppose that $(\mathcal{A}_\alpha)_{\alpha<\omega_1}$ is an ω_1-chain, and φ is a sentence of $\mathcal{L}_{\omega_1,\omega}$ such that $\mathcal{A}_{\omega_1} \models \varphi$. Then, the set $\{\alpha < \omega_1 | A_\alpha \models \varphi\}$ is cofinal in ω_1.*

Proof. Let Φ be the set of subformulas of φ and their negations. Note that Φ is countable. For any $\psi \in \Phi$, $\overline{a} \subset A_{\omega_1}$ such that $\mathcal{A}_{\omega_1} \models \exists \overline{y} \psi(\overline{a}, \overline{y})$, let $\overline{f}_\psi(\overline{a})$ be a tuple of length $\overline{y}$ such that $\mathcal{A}_{\omega_1} \models \psi(\overline{a}, \overline{f}_\psi(\overline{a}))$. For convenience, we think of the $\overline{f}_\psi$ as partial functions on $\mathcal{A}_{\omega_1}$.

Suppose $\beta < \omega_1$. We will show that there is an α such that $\beta \leq \alpha < \omega_1$ and $\mathcal{A}_\alpha \models \varphi$. For $\gamma \geq \beta$, let F_γ be the set of elements of tuples $\overline{f}_\psi(\overline{a})$ for $\overline{a} \subset A_\gamma$, and $\psi \in \Phi$. Because $\mathcal{A}_\gamma$ and Φ are countable, F_γ is countable, so $\overline{A_\gamma} = F_\gamma \cup A_\gamma$ is countable. The set $\{\alpha < \omega_1 | \overline{A_\gamma} \cap (\mathcal{A}_{\alpha+1} \setminus A_\alpha) \neq \emptyset\}$ is then countable, so is not cofinal in ω_1. Therefore, there is some $\gamma' < \omega_1$ which is larger than any element of this set, in which case $\mathcal{A}_{\gamma'} \supset \overline{A_\gamma}$.

Let $\beta_0 = \beta$, and $\beta_{n+1} = \beta_n'$. Let $\delta = \sup\{\beta_n | n < \omega\}$. Then, $\delta < \omega_1$. We claim that $\mathcal{A}_\delta = \bigcup_n A_{\beta_n}$. One of the following is true of the sequence (β_n).

1. For some n, $\beta_n = \beta_{n+1}$.

2. $\beta_n < \beta_{n+1}$ for each n.

In case (1), $\beta_N = \beta_n$ for all $N \geq n$, so $\delta = \beta_n$, and $\mathcal{A}_\delta = A_{\beta_n} = \bigcup_n A_{\beta_n}$. In case (2), δ is a limit ordinal, so $\mathcal{A}_\delta = \bigcup_{\alpha<\delta} A_\alpha = \bigcup_n A_{\beta_n}$. We now claim that $\mathcal{A}_\delta \supset F_\delta$. Suppose $\overline{a} \subset A_\delta$, and $\psi \in \Phi$, such that $\mathcal{A}_{\omega_1} \models \exists \overline{y} \psi(\overline{a}, \overline{y})$. Then, $a \subset A_{\beta_n}$ for some n, so $\overline{f}_\psi(\overline{a}) \in F_{\beta_n} \subset A_{\beta_{n+1}} \subset A_\delta$.

We will now show that for any $\theta \in \Phi$ and $\overline{a} \in A_\delta$, $\mathcal{A}_{\omega_1} \models \theta(\overline{a})$ if and only if $\mathcal{A}_\delta \models \theta(\overline{a})$, by induction on the complexity of θ. If θ is quantifier free, this is true because $\mathcal{A}_\delta \subset A_{\omega_1}$. If $\theta = \neg\eta$ and this is true for η, it follows for θ because $\mathcal{A} \models \theta(\overline{a})$ if and only if $\mathcal{A} \not\models \eta(\overline{a})$. Suppose that this is true for θ_i for $i \in I$, where I is countable. If $\theta = \bigwedge\!\!\!\bigwedge_{i\in I} \theta_i$, then $\mathcal{A}_{\omega_1} \models \theta(\overline{a})$ if and only if $\mathcal{A}_{\omega_1} \models \theta_i(\overline{a})$ for each $i \in I$, if and only if $\mathcal{A}_\delta \models \theta_i(\overline{a})$ for each i, if and only if $\mathcal{A}_\delta \models \theta(\overline{a})$. If $\theta = \bigvee\!\!\!\bigvee_{i\in I} \theta_i$, then $\neg\theta$ is equivalent to $\bigwedge\!\!\!\bigwedge_{i\in I} \neg\theta_i$, in which case this follows from the previous two cases.

Suppose this is true for η. If $\theta = \exists \overline{y} \eta(\overline{y})$ and $\mathcal{A}_{\omega_1} \models \theta(\overline{a})$, then $\mathcal{A}_\delta \models \theta(\overline{a})$ because $\overline{f}_\eta(\overline{a}) \in A_\delta$. If $\mathcal{A}_\delta \models \theta(\overline{a})$, then $\mathcal{A}_{\omega_1} \models \theta(\overline{a})$ because $\mathcal{A}_\delta \subset A_{\omega_1}$. If $\theta = \forall \overline{y} \eta(\overline{y})$, then $\neg\theta$ is equivalent to $\exists \overline{y} \neg\eta(\overline{y})$, so this follows from the previous case and the case of negation. Φ is well founded when ordered by containment of subformulas, so by induction, this is true for any $\theta \in \Phi$.

Because $\varphi \in \Phi$ and $\mathcal{A}_{\omega_1} \models \varphi$, we conclude that $\mathcal{A}_\delta \models \varphi$. Moreover, $\delta \geq \beta$. This completes the proof that $\{\alpha < \omega_1 | A_\alpha \models \varphi\}$ is cofinal.

$\square$

Theorem 3.8. *Suppose that φ is a sentence of $\mathcal{L}_{\omega_1,\omega}$ and $(\mathcal{A}_\alpha)_{\alpha<\omega_1}$ is an ω_1-chain such that for all sufficiently large α, $\mathcal{A}_\alpha \models \varphi$. Then $\mathcal{A}_{\omega_1} \models \varphi$.*

Proof. Suppose $\mathcal{A}_{\omega_1} \models \neg\varphi$. By the previous theorem, $\{\alpha < \omega_1 | A_\alpha \models \neg\varphi\}$ is cofinal, contradicting our hypothesis. $\square$

We can apply this theorem to prove the following.

Theorem 3.9. *Suppose X is a countable structure. Either the $\mathcal{L}_{\omega,\omega}$ theory of X is not countably categorical, or the $\mathcal{L}_{\omega_1,\omega}$ theory of X has an uncountable model.*

Proof. Suppose that the $\mathcal{L}_{\omega,\omega}$ theory of X, $\mathrm{Th}(X)$ is countably categorical. We will first show that X has a proper elementary embedding into itself. By the upward Löwenheim-Skolem theorem, X has an uncountable elementary extension $Y \succ X$. Because Y is uncountable, we can choose some $y \in Y \setminus X$. Applying the downward Löwenheim-Skolem theorem, Y has an elementary substructure X' containing $X \cup \{y\}$. For any finitary formula φ and tuple $\overline{a} \subset X$, if $X \models \varphi(\overline{a})$, then because $X \prec Y$, $Y \models \varphi(\overline{a})$, so because $X' \prec Y$, $X' \models \varphi(\overline{a})$. We conclude that $X \prec Y$. Because $y \notin X$, $X \subsetneq Y$. Because $Y \models \mathrm{Th}(X)$, $Y \simeq X$. Composing the elementary embedding $X \to Y$ with the isomorphism $Y \to X$, we obtain an elementary embedding $X \to X$ whose image is a proper substructure.

Using this elementary embedding, we can define by recursion an ω_1-chain of copies of X such that the successor of each entry in the chain is a proper elementary extension. For countable limit ordinal β, if we have defined X_α for each $\alpha < \beta$, we let $X_\beta = \bigcup_{\alpha<\beta} X_\alpha$. This is then a countable elementary extension of X_0, so is a model of $\mathrm{Th}(X)$, and thus an isomorphic copy of X. Taking the union of this chain, we obtain a countable structure Z. By the previous theorem, for any $\mathcal{L}_{\omega_1,\omega}$ sentence φ such that $X \models \varphi$, $Z \models \varphi$. $\square$

The above theorem can also be proved by other methods. Assuming that the $\mathcal{L}_{\omega_1,\omega}$ theory of X has no uncountable models, we can conclude that X has an uncountable elementary extension not satisfying the Scott sentence of X, and apply the downward Löwenheim-Skolem theorem to this, to obtain a countable structure elementarily equivalent to X, but not isomorphic to it. We could also use the Ryll-Nardzewski theorem to show that if X has a countably categorical theory, the Scott sentence of X is a conjunction of finitary sentences, so is true in some uncountable elementary extension of X, which is then back and forth equivalent to X, and so is a model of its $\mathcal{L}_{\omega_1,\omega}$ theory.

3.4 Morleyization: $n = 3$

Using the method of Morleyization, in which one introduces new symbols to the language to represent formulas, we can extend the methods used in the $n = 2$ case to prove the $n = 3$ case.

Theorem 3.10. *Let ψ be an infinitary Π_3 sentence in signature $\mathcal{L}$, of the form*

$$\psi = \bigwedge_i \forall \overline{x} \bigvee_j \exists \overline{y} \bigwedge_k \forall \overline{z} \theta_{i,j,k}(\overline{x}, \overline{y}, \overline{z})$$

where each $\theta_{i,j,k}$ is finitary and quantifier free. Let T be a finitary $\mathcal{L}$ theory, and φ a finitary $\mathcal{L}$ sentence. Suppose φ is equivalent to ψ over all models of T. Then φ is equivalent to some finitary Π_3 $\mathcal{L}$ sentence over T.

Proof. Consider the Morelyization of T^* with respect to the set of finitary universal formulae, in signature $\mathcal{L}^*$. $\mathcal{L}^*$ consists of $\mathcal{L}$ together with a new relation symbol R_ϕ for each universal formula $\phi(\overline{x}) = \forall \overline{y} \theta(\overline{x}, \overline{y})$. This symbol is defined in T^* by the sentence

$$(\forall \overline{x})(R_\phi(\overline{x}) \leftrightarrow (\forall \overline{y})\theta(\overline{x}, \overline{y}))$$

Note that T^* remains finitary. In models of T^*, ψ, and thus φ, is equivalent to the sentence

$$\psi^* = \bigwedge_i \forall \overline{x} \bigvee_j \exists \overline{y} \bigwedge_k R_{i,j,k}(\overline{x}, \overline{y})$$

where $R_{i,j,k}$ corresponds to the universal formula $\forall \overline{z} \theta_{i,j,k}(\overline{x}, \overline{y}, \overline{z})$. Suppose that

$$\mathcal{A}_0^* \subset A_1^* \subset A_2^* \subset \ldots$$

is a chain of models of T^*, such that $\mathcal{A}_n^* \models \varphi$ for each $n < \omega$, and such that $\mathcal{A}_\omega^* = \bigcup_{n<\omega} A_n^*$ is a model of T^*. We will show that $\mathcal{A}_\omega^* \models \varphi$. We first note that for each n, $\mathcal{A}_n^* \models \psi^*$, and that it suffices to show that $\mathcal{A}_\omega^* \models \psi^*$. Fix some i, and $\overline{a} \subset A_\omega^*$. $\overline{a} \subset A_n^*$ for some n. Because $\mathcal{A}_n^* \models \psi^*$, there is a j and a $\overline{b} \subset A_n^*$ such that

$$\mathcal{A}_n^* \models \bigwedge_k R_{i,j,k}(\overline{a}, \overline{b})$$

That is, for each k, $\mathcal{A}_n^* \models R_{i,j,k}(\overline{a}, \overline{b})$, so $\mathcal{A}_\omega^* \models R_{i,j,k}(\overline{a}, \overline{b})$. Therefore,

$$\mathcal{A}_\omega^* \models \bigvee_j \exists \overline{y} \bigwedge_k R_{i,j,k}(\overline{a}, \overline{y})$$

This holds for any $\overline{a}$, and any i, so $\mathcal{A}_\omega^* \models \psi^*$, which in turn implies that $\mathcal{A}_\omega^* \models \varphi$. We can conclude from this property of φ that φ is equivalent in all models of T^* to a finitary Π_2 sentence η^* in the language $\mathcal{L}^*$. Let η be the sentence obtained by replacing each symbol R_ϕ in η^* with the formula ϕ. Then, η is a finitary Π_3 $\mathcal{L}$ sentence, equivalent to φ is all models of T^*. Because any model of T can be extended to a model of T^*, we conclude that φ and η are equivalent in all models of T. $\square$

Theorem 3.11. *Let $\psi(\overline{x})$ be an infinitary Π_3 formula, equivalent in all models of a finitary theory T to a finitary formula $\varphi(\overline{x})$. Then, φ is equivalent to a finitary Π_3 formula over T.*

Proof. Let $\overline{c}$ be a tuple of new constants, of length $\overline{x}$, and apply the previous theorem to the sentences $\psi(\overline{c})$ and $\varphi(\overline{c})$. $\square$

3.5 Digression: Consistency Properties and Interpolation

Viewing the result we wish to obtain as an interpolation theorem, one might attempt to prove this using methods used to prove other interpolation theorems for infinitary logic. In particular, the method of consistency properties, which has been used to prove Lopez-Escobar's version of the Craig interpolation theorem for $\mathcal{L}_{\omega_1,\omega}$ appears promising. Unfortunately, we can obtain only a very slight reduction in complexity using this method. For a full definition of a consistency property, and its use to prove the Craig interpolation for $\mathcal{L}_{\omega_1,\omega}$ see [Mar16].

For a finite set σ of formulas, we define its complexity $\kappa(\sigma)$ to be the least n such that every element of σ is Π_n. We say that a sentence is $\Pi_n^\dagger$ if it is of the form $\bigvee\!\!\!\bigvee_{i_1} \forall\overline{x_1} \dots \bigvee\!\!\!\bigvee_{i_n} \forall\overline{x_n}\eta_{\overline{i}}(\overline{x})$, where each $\eta_{\overline{i}}$ is a Σ_{n-1} formula. In other words, $\Pi_n^\dagger$ is the smallest class of formulas containing the Σ_{n-1} formulas and closed under universal quantification and infinitary disjunction.

Let $\mathcal{L}$ be a signature, and $\mathcal{L}^*$ be obtained by adding countably many new constants to $\mathcal{L}$. We define Σ to be the set of sets of $\mathcal{L}^*_{\omega_1,\omega}$ sentences $\sigma = \sigma_1 \cup \sigma_2$ with the following properties.

1. σ_1 is finite.
2. σ_2 is a countable set of finitary sentences, containing only finitely many new constants.
3. If $\sigma_1 \vdash \psi$, $\kappa(\sigma_1) = n$, and ψ is $\Pi_n^\dagger$ then $\psi + \sigma_2$ is satisfiable.

Lemma 3.12. *Σ is a consistency property.*

Proof. We will use the following facts. If $\sigma_1 \cup \sigma_2 \in \Sigma$, and $\sigma_1 \vdash \phi$, where $\kappa(\phi) \leq \kappa(\sigma_1)$, then letting $\sigma_1' = \sigma_1 \cup \{\phi\}$, $\sigma_1' \cup \sigma_2 \in \Sigma$. Similarly, if $\sigma_2 \vdash \phi$, where ϕ is finitary, then letting $\sigma_2' = \sigma_2 \cup \{\phi\}$, $\sigma_1 \cup \sigma_2' \in \Sigma$. In both cases, properties (1) and (2) are preserved, and property (3) is preserved because we need only consider ψ that are already considered for $\sigma_1 \cup \sigma_2$.

We will now show that Σ satisfies requirements (1)-(7) to be a consistency property. We let $n = \kappa(\sigma_1)$ throughout.

1. If $\phi, \neg\phi \in \sigma_1$, then $\sigma_1 \vdash c \neq c$ for some c, which is atomic and not satisfiable.

   If $\phi, \neg\phi \in \sigma_2$, then $\sigma_1 \vdash c = c$, and $c = c + \sigma_2$ is not satisfiable.

   If $\phi \in \sigma_1$, $\neg\phi \in \sigma_2$ (or vice versa), then ϕ is finitary, $\kappa(\phi) \leq \kappa(\sigma_1)$ (respectively, $\kappa(\neg\phi) \leq \kappa(\sigma_1)$), so letting $\psi = \phi$ (respectively $\psi = \neg\phi$), we see that $\sigma_1 \cup \sigma_2 \notin \Sigma$.

2. This follows from the facts, as $\neg\phi$ is equivalent to $\sim \phi$, is finitary if and only if $\sim \phi$ is, and is equal to it in complexity.

3. This follows from the facts, as $\bigwedge\!\!\!\bigwedge_i \phi_i \vdash \phi_j$ for each j, ϕ_j is finitary if the conjunction is, and ϕ_j is no more complex than the conjunction.

4. Suppose $\bigvee_i \phi_i \in \sigma_2$. This disjunction is finitary, because σ_2 consists of finitary sentences. Suppose that for each i, $\sigma_1 \cup (\sigma_2 \cup \{\phi_i\}) \notin \Sigma$. Then, for each i, there is a $\Pi_n^\dagger$ ψ_i such that $\sigma_1 \vdash \psi_i$, and $\psi_i + \sigma_2 + \phi_i$ is not satisfiable. Let $\psi = \bigwedge_i \psi_i$. Because this is a finitary conjunction, ψ is $\Pi_n^\dagger$, and $\psi + \sigma_2$ is not satisfiable. This implies that $\sigma_1 \cup \sigma_2 \notin \Sigma$.

   Suppose now that $\bigvee\!\!\!\bigvee_i \phi_i \in \sigma_1$. Then, each ϕ_i is at most Σ_{n-1}. If for each i, $(\sigma_1 \cup \{\phi_i\}) \cup \sigma_2 \notin \Sigma$, there is a $\Pi_n^\dagger$ ψ_i such that $\sigma_1 + \phi_i \vdash \psi_i$ and $\psi_i + \sigma_2$ is not satisfiable. Let $\psi = \bigvee\!\!\!\bigvee_i \psi_i$. Then, ψ is $\Pi_n^\dagger$, $\sigma_1 \vdash \psi$, and $\psi + \sigma_2$ is not satisfiable. This implies that $\sigma_1 \cup \sigma_2 \notin \Sigma$.

   We conclude that property (4) holds by contrapositive.

5. This follows from the facts, as $(\forall x)\varphi(x) \vdash \varphi(c)$ for each c, $\varphi(c)$ is less complex than $(\forall x)\varphi(x)$, and one is finitary if and only if the other is.

6. Suppose $(\exists x)\phi(x) \in \sigma_1$. Then, $\phi(x)$ is at most Σ_{n-1}. Let c be a constant not appearing in σ, and suppose $\sigma_1 + \phi(c) \vdash \theta_1(c)$, for $\theta(c)$ a $\Pi_n^\dagger$ formula. Because c does not occur in σ_1, $\sigma_1 \vdash (\forall x)(\phi(x) \to \theta(x))$, which is $\Pi_n^\dagger$. Furthermore, $(\exists x)\phi(x)$ is $\Pi_n^\dagger$, so

$$\psi = (\exists x)\phi(x) \wedge (\forall x)(\phi(x) \to \theta(x))$$

is $\Pi_n^\dagger$. $\sigma_1 \vdash \psi$, so $\psi + \sigma_2$ is satisfiable. $\psi \vdash (\exists x)\theta(x)$, so $(\exists x)\theta(x) + \sigma_2$ is satisfiable. Because c does not occur in σ_2, this implies that $\theta(c) + \sigma_2$ is satisfiable. Therefore, $(\sigma_1 \cup \{\phi(c)\}) \cup \sigma_2 \in \Sigma$.

Suppose that $(\exists x)\phi(x) \in \sigma_2$. Again, let c be a constant not occurring in σ. Suppose $\sigma_1 \vdash \theta(c)$, a $\Pi_n^\dagger$ sentence. Because c does not occur in σ_1, $\sigma_1 \vdash (\forall x)\theta(x)$, which is a $\Pi_n^\dagger$ sentence. Therefore, $(\forall x)\theta(x) + \sigma_2$ is satisfiable, in a model M. Interpret c as a witness to $(\exists x)\phi(x)$ in M. Then, $M \models \theta(c) + \sigma_2 + \phi(c)$. We conclude that $\sigma_1 \cup (\sigma_2 \cup \{\phi(c)\}) \in \Sigma$.

7. (a) This follows from the facts, as $d = c$ is atomic and a consequence of $c = d$.
   (b) This follows from the facts, as $\phi(c)$ follows from $c = t$ and $\phi(t)$, is no more complex than $\phi(t)$, and is finitary if $\phi(t)$ is.
   (c) Let c be a constant not occurring in σ. Suppose $\sigma_1 + c = t \vdash \theta(c)$, a $\Pi_n^\dagger$ sentence. Then, $\sigma_1 \vdash \theta(t)$, which is $\Pi_n^\dagger$, so $\theta(t) + \sigma_2$ is satisfiable, in a model M. Interpreting c as t^M, we obtain a model of $\theta(c) + \sigma_2$. Therefore, $(\sigma_1 \cup \{c = t\}) \cup \sigma_2 \in \Sigma$.

□

Theorem 3.13 ($\Pi_n^\dagger$ Interpolation)**.** *Suppose that ψ is a Π_n sentence, and φ is a finitary sentence such that ψ is equivalent to φ in all models of T, a finitary theory. Then, ψ is equivalent to a $\Pi_n^\dagger$ sentence in all models of T.*

Proof. There is no model of $T + \psi + \neg\varphi$, so letting $\sigma_1 = \{\psi\}$, and $\sigma_2 = T \cup \{\neg\varphi\}$, $\sigma_1 \cup \sigma_2 \notin \Sigma$. Therefore, there is a $\Pi_n^\dagger$ sentence $\theta(\overline{c})$, for $\overline{c}$ a sequence of new constants, such that $\psi \vdash \theta(\overline{c})$, and $\theta(\overline{c}) + T + \neg\varphi$ is not satisfiable. Then, in models of T, $\psi \vdash (\forall\overline{x})\theta(\overline{x}) \vdash (\exists\overline{x})\theta(\overline{x}) \vdash \varphi \vdash \psi$, so $(\forall\overline{x})\theta(\overline{x})$ is a $\Pi_n^\dagger$ sentence equivalent to ψ. □

3.6 n-Elementary Extensions, $n = 4$

In this case, we will prove the contrapositive of the $n = 4$ case, which will involve a useful semantic test for Π_n formulas more generally. We say that $\mathcal{A}$ is an n-elementary substructure of $\mathcal{B}$, and write $\mathcal{A} \prec_n \mathcal{B}$ if $\mathcal{A} \subset \mathcal{B}$ and for any $\forall_n$ formula φ, and elements $\overline{a}$ of $\mathcal{A}$, if $\mathcal{A} \models \varphi(\overline{a})$, then $\mathcal{B} \models \varphi(\overline{a})$. We let $E_{\forall_n}(\mathcal{A})$ be the set of $\forall_n$ sentences in the elementary diagram of $\mathcal{A}$. Then, $\mathcal{A} \prec_n \mathcal{B}$ if and only if $\mathcal{B} \models E_{\forall_n}(\overline{A})$, interpreting constants naming elements of $\mathcal{A}$ as the same elements of $\mathcal{B}$. Any model of $E_{\forall_n}(\mathcal{A})$ can be replaced with an isomorphic n-elementary extension of $\mathcal{A}$, replacing the interpretations of constants naming elements of $\mathcal{A}$ with those elements. Under this definition, $\mathcal{A} \prec_0 \mathcal{B}$ if and only if $\mathcal{A} \subset \mathcal{B}$, and $\mathcal{A} \prec \mathcal{B}$ if and only if $\mathcal{A} \prec_n \mathcal{B}$ for all n. Another way to interpret this definition is that an n-elementary extension is an extension of structures, after Moreleyizing with respect to finitary Π_n formulas.

Let $\mathcal{L}$ be countable. Suppose $\psi = \bigwedge\!\!\!\bigwedge_i \forall\overline{u} \bigvee\!\!\!\bigvee_j \exists\overline{v} \bigwedge\!\!\!\bigwedge_k \forall\overline{w} \bigvee\!\!\!\bigvee_\ell \exists\overline{x}\theta_{i,j,k,\ell}(\overline{uvwx})$ is an $\mathcal{L}_{\omega_1,\omega}$ sentence and φ is a $\mathcal{L}_{\omega,\omega}$ sentence not equivalent to any finitary Π_4 $\mathcal{L}$ sentence over a finitary $\mathcal{L}$ theory T. We will attempt to construct a model of T showing that $\psi \not\leftrightarrow_T \varphi$.

Let Γ be the set of Π_4 consequences of $\varphi + T$ in $\mathcal{L}_{\omega,\omega}$. Because φ is not equivalent to any finitary Π_4 sentence over T, $\Gamma + T \not\vdash \varphi$. Let $\mathcal{A}_0$ be a countable model of $\Gamma + T + \neg\varphi$. We now claim that $\varphi + T + E_{\Pi_3}(\mathcal{A})$ is satisfiable. If not, then for some $\phi_1(\overline{c}), \ldots, \phi_n(\overline{c}) \in E_{\Pi_3}(\mathcal{A})$, $\varphi + T \vdash \neg\exists\overline{x} \bigwedge_{i\leq n} \phi_i(\overline{x})$, which is Π_4, so is an element of Γ, contradicting the fact that $\mathcal{A}_0 \models \Gamma$.

Let $\mathcal{B}_0^0 \models E_{\Pi_3}(\mathcal{A}) + T + \varphi$, so $\mathcal{A}_0 \prec_3 \mathcal{B}_0^0$. Pick an enumeration $\{(i_0^0, \overline{a}_0^0), (i_1^0, \overline{a}_1^0), \dots\}$ of pairs consisting of an i and a tuple of length $\overline{u}$ in $\mathcal{A}_0$. We start with $\overline{a}_0^0 \in \mathcal{B}_0^0$. Consider the tuples $\overline{b} \in \mathcal{B}_0^0$ of length $\overline{v}$. We say a pair $(j, \overline{b})$ is breakable if there is a $k_{j,\overline{b}}$ such that for each L there is a $\overline{c}_L$ of length $\overline{w}$ such that

$$\mathcal{B}_0^0 \models \bigwedge_{\ell \leq L} \forall \overline{x} \neg \theta_{i_0^0, j, k_{j,\overline{b}}, \ell}(\overline{a}_0^0 \overline{b} \overline{c}_L \overline{x})$$

Otherwise, we say that it is unbreakable.

Suppose all $(j, \overline{b})$ are breakable. Then,

$$E(\mathcal{B}_0^0) + \{\forall \overline{x} \neg \theta_{i_0^0, j, k_{j,\overline{b}}, \ell}(\overline{a}_0^0 \overline{b} \overline{c}_{j,\overline{b},k} \overline{x}) | (j, \overline{b}), \ell\}$$

is finitely satisifiable, where $\overline{c}_{j,\overline{b},k}$ are new constants. This set is then satisifiable, in a model $\mathcal{B}_0^1$. Then, $\mathcal{B}_0^1 \succ \mathcal{B}_0^0$, and for each $(j, \overline{b})$,

$$\mathcal{B}_0^1 \models \exists \overline{w} \bigwedge\!\!\!\bigwedge_{\ell} \forall \overline{x} \neg \theta_{i_0^0, j, k_{j,\overline{b}}, \ell}(\overline{a}_0^0 \overline{b} \overline{w} \overline{x})$$

Moreover, any elementary extension of $\mathcal{B}_0^1$ satisfies this, using the same witness for $\overline{w}$. We can then iterate this to construct structures $\mathcal{B}_0^n \prec \mathcal{B}_0^{n+1}$, so long as each $(j, \overline{b})$ with $\overline{b} \in \mathcal{B}_0^n$ is breakable.

Suppose conversely that we get to a $\mathcal{B}_0^n$ (maybe $\mathcal{B}_0^0$) with an unbreakable $(j, \overline{b})$. Then, for every k, there is an L_k such that for every $\overline{c}$ of length $\overline{w}$,

$$\mathcal{B}_0^n \models \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{b} \overline{c} \overline{x})$$

In other words,

$$\mathcal{B}_0^n \models \bigwedge\!\!\!\bigwedge_{k} \forall \overline{w} \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{b} \overline{w} \overline{x})$$

We conclude that for each K,

$$\mathcal{B}_0^n \models \exists \overline{v} \bigwedge_{k \leq K} \forall w \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{v} \overline{w} \overline{x})$$

This is a finitary Σ_3 statement, so because $\mathcal{A}_0 \prec_3 \mathcal{B}_0^n$

$$\mathcal{A}_0 \models \exists \overline{v} \bigwedge_{k \leq K} \forall \overline{w} \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{v} \overline{w} \overline{x})$$

for each K.

We now have two cases. Either

$$\mathcal{A}_0 \models \exists \overline{v} \bigwedge_{k \leq K} \forall w \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{v} \overline{w} \overline{x})$$

for each K, or we construct an infinite elementary chain

$$\mathcal{B}_0^0 \prec \mathcal{B}_0^1 \prec \mathcal{B}_0^2 \prec \dots$$

We will deal with the second case first.

Let $\mathcal{B} = \bigcup_{n<\omega} \mathcal{B}_0^n$. Then, $\mathcal{B}_0^n \prec B$ for each n. This implies that $\mathcal{B} \models T + \varphi$. We will show that $\mathcal{B} \models \neg\psi$. This is equivalent to $\bigvee\!\!\!\bigvee_i \exists \overline{u} \bigwedge\!\!\!\bigwedge_j \forall \overline{v} \bigvee\!\!\!\bigvee_k \exists \overline{w} \bigwedge\!\!\!\bigwedge_\ell \forall \overline{x} \neg \theta_{i,j,k,\ell}(\overline{uvwx})$. We will take i_0^0, $\overline{a}_0^0$ as witnesses for i and $\overline{u}$. It then suffices to show that

$$\mathcal{B} \models \bigwedge\!\!\!\bigwedge_{j} \forall \overline{v} \bigvee\!\!\!\bigvee_{k} \exists \overline{w} \bigwedge\!\!\!\bigwedge_{l} \forall \overline{x} \neg \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{v} \overline{w} \overline{x})$$

We will consider each pair $(j, \bar{b})$, for $\bar{b} \in B$ of length $\overline{v}$, one at a time. This $\bar{b}$ shows up in some $\mathcal{B}_0^n$, and the pair $(j, \bar{b})$ is "broken" in the construction of $\mathcal{B}_0^{n+1}$. In doing so, we ensure that because $\mathcal{B}$ is an elementary extension of $\mathcal{B}_0^{n+1}$ it satisfies

$$\mathcal{B} \models \exists \overline{w} \bigwedge\!\!\!\bigwedge_{\ell} \forall \overline{x} \neg \theta_{i_0^0, j, k_{j,\bar{b}}, \ell}(\overline{a}_0^0 \bar{b} \overline{w} \overline{x})$$

Thus,

$$\mathcal{B} \models \bigvee\!\!\!\bigvee_{k} \exists \overline{w} \bigwedge\!\!\!\bigwedge_{\ell} \forall \overline{x} \neg \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \bar{b} \overline{w} \overline{x})$$

This holds for all $(j, \bar{b})$, so we conclude that $\mathcal{B} \models \neg\psi$.

We now consider the first case. Let $\overline{d}$ be new constants of length $\overline{v}$.

$$E(\mathcal{A}_0) + \{ \bigwedge_{k \leq K} \forall \overline{w} \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{d} \overline{w} \overline{x}) | K \}$$

is finitely satisfiable, so satisfiable, in a countable model $\mathcal{A}_1$. Then, $\mathcal{A}_1 \succ A_0$, and

$$\mathcal{A}_1 \models \exists \overline{v} \bigwedge\!\!\!\bigwedge_{k} \forall \overline{w} \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i_0^0, j, k, \ell}(\overline{a}_0^0 \overline{v} \overline{w} \overline{x})$$

witnessed by $\overline{d}^{A_1}$. As with $\mathcal{A}_0$, $\mathcal{A}_1 \models \Gamma + T$, so there is a $\mathcal{B}_1^0$ such that $\mathcal{A}_1 \prec_3 \mathcal{B}_1^0$ and $\mathcal{B}_1^0 \models \varphi + T$.

We can then iterate the entire process, subject to the following modification. Let $f : \mathbb{N} \to \mathbb{N}^2$ be a bijection, with components g and h, such that $g(n) \leq n$ for all n, and $f(0) = (0, 0)$. When we construct a $\mathcal{A}_n$, choose an enumeration $\{(i_0^n, \overline{a}_0^n), (i_1^n, \overline{a}_1^n), \dots\}$ of pairs consisting of an i and a tuple of length $\overline{u}$ in $\mathcal{A}_n$. After constructing $\mathcal{A}_n$, we use $(i_{h(n)}^{g(n)}, \overline{a}_{h(n)}^{g(n)})$ in place of $(i_0^0, \overline{a}_0^0)$.

We again have two cases. Either at some point, we construct a $\mathcal{B}$ such that $\mathcal{B} \models T + \varphi + \neg\psi$, or we construct an infinite elementary chain

$$\mathcal{A}_0 \prec A_1 \prec A_2 \prec \dots$$

Let $\mathcal{A} = \bigcup_{n < \omega} A_n$. For each n, $\mathcal{A}_n \models T + \neg\varphi$, so $\mathcal{A} \models T + \neg\varphi$. We will show that $\mathcal{A} \models \psi$. Fix some i, $\overline{a} \in A$. The pair $(i, \overline{a})$ shows up as $(i_{h(N)}^{g(N)}, \overline{a}_{h(N)}^{g(N)})$ at some step N, after we have constructed $\mathcal{A}_N$. We construct $\mathcal{A}_{N+1}$ so that

$$\mathcal{A}_{N+1} \models \exists \overline{v} \bigwedge\!\!\!\bigwedge_{k} \forall \overline{w} \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i,j,k,\ell}(\overline{avwx})$$

Let $\bar{b}$ be a witness to this. Then, for each k

$$\mathcal{A}_{N+1} \models \forall \overline{w} \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i,j,k,\ell}(\overline{abwx})$$

Because $\mathcal{A} \succ A_{N+1}$, for each k,

$$\mathcal{A} \models \forall \overline{w} \bigvee_{\ell \leq L_k} \exists \overline{x} \theta_{i,j,k,\ell}(\overline{abwx})$$

which in turn implies that, for each k,

$$\mathcal{A} \models \forall \overline{w} \bigvee\!\!\!\bigvee_{\ell} \exists \overline{x} \theta_{i,j,k,\ell}(\overline{abwx})$$

Consequently,

$$\mathcal{A} \models \exists \overline{v} \bigwedge\!\!\!\bigwedge_{k} \forall \overline{w} \bigvee\!\!\!\bigvee_{\ell} \exists \overline{x} \theta_{i,j,k,\ell}(\overline{avwx})$$

This holds for any $i, \overline{a}$, so $\mathcal{A} \models \psi$.

We conclude that there is either a $\mathcal{B} \models T + \varphi + \neg\psi$ or an $\mathcal{A} \models T + \neg\varphi + \psi$, so $\psi \not\leftrightarrow_T \varphi$. By contrapositive, we have the following.

Theorem 3.14. *Suppose ψ is a Π_4 sentence, equivalent to a finitary sentence φ over a finitary theory T in a countable language. Then, ψ and φ are equivalent to a finitary Π_4 sentence over T*

3.7 The $n = 5$ Case

We can, with a more elaborate construction, prove the $n = 5$ case by a similar method to the $n = 4$ case.

We will first state and prove the general form of the semantic test used in the $n = 4$ case.

Theorem 3.15. *Let φ be a finitary formula, and T be a finitary theory. The following are equivalent.*

1. *φ is equivalent to a finitary $\forall_{n+1}$ formula over T.*
2. *If $\mathcal{A} \prec_n \mathcal{B}$ are models of T, and $\overline{a} \in \mathcal{A}$, then either $\mathcal{A} \models \varphi(\overline{a})$, or $\mathcal{B} \models \neg\varphi(\overline{a})$. That is, φ is preserved under n-elementary submodels of models of T.*

Proof. Suppose that φ is equivalent to a finitary $\forall_{n+1}$ formula η over T. It suffices to show that η has property (2). We may assume that η is in prenex normal form, so that $\eta(\overline{x}) = \forall\overline{y}\theta(\overline{xy})$, for θ a $\exists_n$ formula. Suppose $\mathcal{B} \models \eta(\overline{a})$. Then, for any $\overline{b} \in \mathcal{A}$, $\overline{b} \in \mathcal{B}$, so $\mathcal{B} \models \theta(\overline{ab})$. Because $\mathcal{A} \prec_n \mathcal{B}$ and θ is $\exists_n$, $\mathcal{A} \models \theta(\overline{ab})$. We conclude that $\mathcal{A} \models \eta(\overline{a})$. Therefore, either $\mathcal{A} \models \eta(\overline{a})$ or $\mathcal{B} \models \neg\eta(\overline{a})$.

Suppose now that φ is not equivalent to any $\forall_{n+1}$ formula over T. Let $\overline{c}$ be new constants, and Γ be the set of $\forall_{n+1}$ consequences of $T + \varphi(\overline{c})$ in the expanded language. We will show that $T + \Gamma \not\vdash \varphi(\overline{c})$. Otherwise, the compactness theorem implies that for some finite $S \subset \Gamma$, $T + S \vdash \varphi(\overline{c})$. Let $\eta(\overline{c}) = \bigwedge_{\theta \in S} \theta$. Then, $\eta(\overline{c})$ is $\forall_{n+1}$, and $T \vdash \eta(\overline{c}) \leftrightarrow \varphi(\overline{c})$. The constants $\overline{c}$ do not occur in T, so η and φ are equivalent over T, contradicting our assumption. Because $T + \Gamma \not\vdash \varphi$, there is some $\mathcal{A} \models T + \Gamma + \neg\varphi(\overline{c})$. We will show that $T + E_{\forall_n}(\mathcal{A}) + \varphi(\overline{c})$ is satisfiable. If not, then for some finite $R \subset E_{\forall_n}(\mathcal{A})$, $T + R + \varphi(\overline{c})$ is not satisfiable. Let $\lambda(\overline{d}) = \bigwedge_{\rho \in R} \rho$, where $\overline{d}$ are constants naming elements of $\mathcal{A}$. Then, because the constants $\lceil$ do not occur in T or $\varphi(\overline{c})$, $T + \varphi(\overline{c}) \vdash \forall\overline{x}\neg\lambda(\overline{x})$. This is a $\forall_{n+1}$ sentence, so is an element of Γ. However, $\mathcal{A} \models \exists\overline{x}\lambda(\overline{x})$, which contradicts the fact that $\mathcal{A} \models \Gamma$. We conclude that $T + E_{\forall_n}(\mathcal{A}) + \varphi(\overline{c})$ is satisfiable. Let $\mathcal{B} \models T + E_{\forall_n}(\mathcal{A}) + \varphi(\overline{c})$. Then, $\mathcal{A} \prec_n \mathcal{B}$. Let $\overline{a}$ be the interpretation of $\overline{c}$ in $\mathcal{A}$. Then, $\mathcal{A} \models \neg\varphi(\overline{a})$, and $\mathcal{B} \models \varphi(\overline{a})$, so φ does not satisfy (2). $\square$

We will also need the following amalgamation results. Here $n \geq 0$.

Lemma 3.16. *Suppose $\mathcal{A} \prec_n \mathcal{B}$ and $\mathcal{A} \prec \mathcal{A}'$. There is a $\mathcal{B}'$ such that $\mathcal{B}' \succ \mathcal{B}$ and $\mathcal{B}' \succ_n \mathcal{B}$.*

Proof. It suffices to show that $E(\mathcal{B}) + E_{\forall_n}(\mathcal{A}')$ is satisifiable, identifying constants naming elements of $\mathcal{A}$ in both sets. Taking a conjunction of a finite subset of $E_{\forall_n}(\mathcal{A}')$, it suffices to show that any sentence $\alpha(\overline{a}, \overline{a}') \in E_{\forall_n}(\mathcal{A}')$, with $\overline{a}$ naming elements of $\mathcal{A}$, can be satisfied in $\mathcal{B}$, by choosing interpretations for $\overline{a}'$. In this case, $\mathcal{A}' \models \exists\overline{y}\alpha(\overline{a}, \overline{y})$, so because $\mathcal{A} \prec \mathcal{A}'$, $\mathcal{A} \models \exists\overline{y}\alpha(\overline{a}, \overline{y})$. This is $\exists_{n+1}$, so because $\mathcal{A} \prec_n \mathcal{B}$, $\mathcal{B} \models \exists\overline{y}\alpha(\overline{a}, \overline{y})$. We conclude that $E(\mathcal{B}) + E_{\forall_n}(\mathcal{A}')$ is satisfiable. A model of this is the required $\mathcal{B}'$. $\square$

Lemma 3.17. *If $\mathcal{A} \prec_{n+1} \mathcal{B}$, then there is a $\mathcal{C} \succ_n \mathcal{B}$ with $\mathcal{C} \succ \mathcal{A}$.*

Proof. It suffices to show that $E(\mathcal{A}) + E_{\forall_n}(\mathcal{B})$ is satisfiable, identifying constants naming elements of $\mathcal{A}$ in both sets. Taking a conjunction of a finite subset of $E_{\forall_n}(\mathcal{B})$, it suffices to show that any sentence $\alpha(\overline{a}, \overline{b}) \in E_{\forall_n}(\mathcal{B})$ where $\overline{a}$ names elements of $\mathcal{A}$ can be satisfied in $\mathcal{A}$, by choosing interpretations for $\overline{b}$. In this case, $\mathcal{B} \models \exists\overline{y}\alpha(\overline{a}, \overline{y})$, which is $\exists_{n+1}$, so because $\mathcal{A} \prec_{n+1} \mathcal{B}$, $\mathcal{A} \models \exists\overline{y}\alpha(\overline{a}, \overline{y})$. We conclude that $E(\mathcal{A}) + E_{\forall_n}(\mathcal{B})$ is satisfiable. A model of this is the required $\mathcal{C}$. $\square$

We now turn to the proof of the $n = 5$ case.

Suppose

$$\psi = \bigwedge\!\!\!\bigwedge_i \forall\overline{u} \bigvee\!\!\!\bigvee_j \exists\overline{v} \bigwedge\!\!\!\bigwedge_k \forall\overline{w} \bigvee\!\!\!\bigvee_\ell \exists\overline{x} \bigwedge\!\!\!\bigwedge_m \forall\overline{y}\theta_{i,j,k,\ell,m}(\overline{uvwxy})$$

and that $\mathcal{A}_0 \prec_4 \mathcal{B}_0$. We will construct either an $\mathcal{A}' \succ A_0$ with $\mathcal{A}' \models \psi$ or a $\mathcal{B}' \succ \mathcal{B}_0$ with $\mathcal{B}' \models \neg\psi$.

By the first amalgamation result, there is a $\mathcal{C}_0 \succ_3 \mathcal{B}_0$ with $C_0 \succ A_0$. Fix $i_0, \overline{a}_0 \in \mathcal{A}_0$ and $j_0, \overline{b}_0 \in \mathcal{B}_0$. We say that a pair $(k, \overline{c})$ with $\overline{c} \in \mathcal{C}_0$ is breakable if there is an $\ell_{k,\overline{c}}$ such that for all M there is a $\overline{d}_M \in \mathcal{C}_0$ such that

$$\mathcal{C}_0 \models \bigwedge_{m \leq M} \forall \overline{y} \theta_{i_0,j_0,k,\ell_{k,\overline{c}},m}(\overline{a_0 b_0 c d_M y})$$

and unbreakable otherwise. Suppose all $(k, \overline{c})$ with $\overline{c} \in \mathcal{C}_0$ are breakable. Let $\overline{d}_{k,\overline{c}}$ be new constants. Then

$$E(\mathcal{C}_0) + \left\{ \forall \overline{y} \theta_{i_0,j_0,k,\ell_{k,\overline{c}},m}(\overline{a_0 b_0 c d_{k,\overline{c}} y}) \middle| (\overline{k}, \overline{c}), m \right\}$$

is finitely satisfiable, so satisfiable, with a model $\mathcal{C}_1$. $\mathcal{C}_1 \succ \mathcal{C}_0$, and for each $k, \overline{c} \in \mathcal{C}_0$,

$$\mathcal{C}_1 \models \bigvee\!\!\!\bigvee_{\ell} \exists \overline{x} \bigwedge\!\!\!\bigwedge_{m} \forall \overline{y} \theta_{i_0,j_0,k,\ell,m}(\overline{a_0 b_0 c x y})$$

Moreover, this is true of any elementary extension of $\mathcal{C}_1$, using the same witnesses. We can iterate this process to construct an elementary chain $\mathcal{C}_0 \prec \mathcal{C}_1 \prec \ldots$ so long as all pairs $(k, \overline{c})$ with $\overline{c} \in \mathcal{C}_n$ are breakable for each n. Suppose conversely that for some n, there is an unbreakable pair $(k, \overline{c})$ with $\overline{c} \in \mathcal{C}_n$. Then for each ℓ there is an M_ℓ such that

$$\mathcal{C}_n \models \forall \overline{x} \bigvee_{m \leq M_\ell} \exists \overline{y} \neg \theta_{i_0,j_0,k,\ell,m}(\overline{a_0 b_0 c x y})$$

That is, for each L,

$$\mathcal{C}_n \models \exists \overline{w} \bigwedge_{\ell \leq L} \forall \overline{x} \bigvee_{m \leq M_\ell} \exists \overline{y} \neg \theta_{i_0,j_0,k,\ell,m}(\overline{a_0 b_0 w x y})$$

This is a finitary Σ_3 formula, so is true in $\mathcal{B}_0$ as well, because $\mathcal{B}_0 \prec_3 C_n$. Let $\overline{c}_*$ be new constants. Then

$$E(\mathcal{B}_0) + \left\{ \forall \overline{x} \bigvee_{m \leq M_\ell} \exists \overline{y} \neg \theta_{i_0,j_0,k,\ell,m}(\overline{a_0 b_0 c_* x y}) \middle| \ell \right\}$$

is finitely satisifiable, so satisifiable, with a model $\mathcal{B}_1$. $\mathcal{B}_1 \succ \mathcal{B}_0$, and

$$\mathcal{B}_1 \models \bigvee\!\!\!\bigvee_{k} \exists \overline{w} \bigwedge\!\!\!\bigwedge_{\ell} \forall \overline{x} \bigvee_{m \leq M_\ell} \exists \overline{y} \neg \theta_{i_0,j_0,k,\ell,m}(\overline{a_0 b_0 w x y})$$

as does any elementary extension of $\mathcal{B}_1$, using the same witnesses. Using the second amalgamation result, we have a new $\mathcal{C}_0' \succ_3 \mathcal{B}_1$, with $\mathcal{C}_0' \succ \mathcal{C}_0$. We can then iterate this process, to construct an elementary chain $\mathcal{B}_0 \prec \mathcal{B}_1 \prec \ldots$ so long as for each $\mathcal{B}_n$, there is an associated $\mathcal{C}_m$ with an unbreakable pair. For each n, we fix a $j_n, \overline{b}_n \in \mathcal{B}_n$. These choices are unconstrained and all structures are countable, so we can choose these so that every such pair is eventually covered.

We then have two cases. Either we construct an infinite chain

$$\mathcal{B}_0 \prec \mathcal{B}_1 \prec \ldots$$

covering all pairs $j, \overline{b}$ or for some $\mathcal{B}_n$, we construct an infinite chain

$$\mathcal{B}_n \prec_3 \mathcal{C}_0 \prec \mathcal{C}_1 \prec \ldots$$

In the first case, let $\mathcal{B} = \bigcup_n \mathcal{B}_n$. Then, $\mathcal{B} \succ \mathcal{B}_0$, and for each j, $\overline{b} \in B$, we have that

$$\mathcal{B} \models \bigvee\!\!\!\bigvee_{k} \exists \overline{w} \bigwedge\!\!\!\bigwedge_{\ell} \forall \overline{x} \bigvee_{m \leq M_\ell} \exists \overline{y} \neg \theta_{i_0,j_0,k,\ell,m}(\overline{a_0 b w x y})$$

so in fact, $\mathcal{B} \models \neg\psi$, witnessed by $i_0, \overline{a}_0$.

In the second case, let $\mathcal{C} = \bigcup_n \mathcal{C}_n$. Then, for each $k, \overline{c} \in \mathcal{C}$,

$$\mathcal{C} \models \bigvee_{\ell} \exists \overline{x} \bigwedge_{m} \forall \overline{y} \theta_{i_0,j_n,k,\ell,m}(\overline{a_0 b_n c x y})$$

so in fact

$$\mathcal{C} \models \bigwedge_{k} \forall \overline{w} \bigvee_{\ell} \exists \overline{x} \bigwedge_{m} \forall \overline{y} \theta_{i_0,j_n,k,\ell,m}(\overline{a_0 b_n w x y})$$

We now have two cases. Either for some elementary extension $\mathcal{D} \succ \mathcal{C}$, there is an unbreakable pair $(k, \overline{d})$ with $\overline{d} \in \mathcal{D}$, or for all $\mathcal{D} \succ \mathcal{C}$, every pair $(k, \overline{d})$ with $\overline{d} \in D$ is breakable. In the first case, using $\mathcal{D} \succ_3 \mathcal{B}_N$ with an unbreakable pair, we can resume the construction of the chain

$$\mathcal{B}_0 \prec \mathcal{B}_1 \prec \cdots \prec \mathcal{B}_N \prec \mathcal{B}_{N+1} \prec \ldots$$

This can be repeated to construct an infinite chain $\mathcal{B}_0 \prec \mathcal{B}_1 \prec \ldots$ as before, or until the second case arises. In the second case, we let $\mathcal{A}_1 = \mathcal{C}$. We say that $(i_0, \overline{a_0})$ is witnessed in $\mathcal{A}_1$ by $(j_n, \overline{b}_n)$. Using the second amalgamation result, we have a $\mathcal{B}_0' \succ \mathcal{B}_0$ such that $\mathcal{B}_0' \succ_4 A_1$.

We can then iterate this process to construct a chain

$$\mathcal{A}_0 \prec A_1 \prec A_2 \prec \ldots$$

so long as at each step, we fail to construct a $\mathcal{B}$. At each step we fix $i_n, \overline{a}_n$. These choices are unconstrained, so because each structure is countable, we can choose them so that all are eventually covered. For each n, $(i_n, \overline{a}_n)$ will be witnessed in $\mathcal{A}_{n+1}$ by some $(j_N, \overline{b}_N)$.

Let $\mathcal{A}_\omega = \bigcup_n A_n$. Then $\mathcal{A}_\omega \succ A_0$. Consider a pair $(i, \overline{a})$ with $\overline{a} \in A$. This shows up as some $(i_n, \overline{a}_n)$. Let it be witnessed by $(j(i, \overline{a}), \overline{b}(i, \overline{a}))$. Because $\mathcal{A}_\omega \succ A_{n+1}$, every pair $(k, \overline{c})$ with $\overline{c} \in A_\omega$ is breakable with respect to $(i, \overline{a})$, $(j(i, \overline{a}), \overline{b}(i, \overline{a}))$. We conclude that, for some indices $\ell(i, \overline{a}, k, \overline{c})$ and new constants $\overline{d}(i, \overline{a}, k, \overline{c})$, the set

$$E(\mathcal{A}_\omega) + \left\{\forall \overline{y} \theta_{i,j(i,\overline{a}),k,\ell(i,\overline{a},k,\overline{c}),m}(\overline{a}\overline{b}(i,\overline{a})\overline{c}\overline{d}(i,\overline{a},k,\overline{c})\overline{y}) \middle| (i,\overline{a}),(k,\overline{c}),m\right\}$$

is finitely satisfiable, so is satisfiable, in a model $\mathcal{A}_{\omega+1} \succ A_\omega$. Then for any $i, \overline{a} \in A_\omega$, for any $k, \overline{c} \in A_\omega$

$$\mathcal{A}_{\omega+1} \models \bigvee_{\ell} \exists \overline{x} \bigwedge_{m} \forall \overline{y} \theta_{i,j(i,\overline{a}),k,\ell,m}(\overline{a}\overline{b}(i,\overline{a})\overline{cxy})$$

This holds in any elementary extension of $\mathcal{A}_{\omega+1}$. $\mathcal{A}_{\omega+1}$ however, in general contains new elements not present in $\mathcal{A}_\omega$.

Using the second amalgamation result, there is a $\mathcal{B}_{\omega+1} \succ_4 A_{\omega+1}$ such that $\mathcal{B}_{\omega+1} \succ \mathcal{B}_0$. Any elementary extension of $\mathcal{B}_{\omega+1}$ is then an elementary extension of $\mathcal{B}_0$, so we can iterate the construction of $\mathcal{A}_{\omega+1}$, starting with $\mathcal{A}_{\omega+1} \prec_4 \mathcal{B}_{\omega+1}$ with the following modification. Any pairs $i, \overline{a}$ already given witnesses in the construction of $\mathcal{A}_{\omega+1}$ keep their old witnesses, only new pairs are given new witnesses. Assuming that at no point we construct a $\mathcal{B} \succ \mathcal{B}_0$ with $\mathcal{B} \models \neg\psi$, we construct an infinite elementary chain

$$\mathcal{A}_0 \prec A_{\omega+1} \prec A_{\omega 2+1} \prec A_{\omega 3+1} \prec \ldots$$

Let $\mathcal{A}_{\omega^2}$ be the limit of this chain. Then, $\mathcal{A}_{\omega^2} \succ A_0$. We will show that $\mathcal{A}_{\omega^2} \models \psi$. Fix $i, \overline{a} \in A_{\omega^2}$. This shows up in some $\mathcal{A}_{\omega n+1}$, and receives a witness $(j(i, \overline{a}), \overline{b}(i, \overline{a}))$ in the construction of $\mathcal{A}_{\omega(n+1)}$. Now fix some $k, \overline{c} \in A_{\omega^2}$. This shows up in some $\mathcal{A}_{\omega n'}$ with $n' > n+1$. In the construction of $\mathcal{A}_{\omega n'+1}$ we ensure that because $\mathcal{A}_{\omega^2}$ is an elementary extension of $\mathcal{A}_{\omega n'+1}$,

$$\mathcal{A}_{\omega^2} \models \bigvee_{\ell} \exists \overline{x} \bigwedge_{m} \forall \overline{y} \theta_{i,j(i,\overline{a}),k,\ell,m}(\overline{a}\overline{b}(i,\overline{a})\overline{cxy})$$

We conclude that $\mathcal{A}_{\omega^2} \models \psi$. The only way the construction of $\mathcal{A}_{\omega^2}$ can fail is if at some point we construct a $\mathcal{B} \succ \mathcal{B}_0$ with $\mathcal{B} \models \neg\psi$.

Theorem 3.18. *If φ is a finitary sentence equivalent over a finitary countable theory T to an infinitary Π_5 sentence ψ, φ is equivalent over T to a finitary Π_5 sentence.*

Proof. Suppose that φ is not equivalent to any Π_5 sentence over T. Then, by Theorem 3.15 we can construct $\mathcal{A}_0 \prec_4 \mathcal{B}_0$, models of T, such that $\mathcal{A}_0 \models \neg\varphi$ and $\mathcal{B}_0 \models \varphi$. The above construction yields either a model of $T + \neg\varphi + \psi$ or $T + \varphi + \neg\psi$, contradicting our assumption. $\square$

In the $n = 5$ case, we check if something can be realized in some elementary extension of a structure, and if so, we realize it. This idea can be isolated and generalized substantially, yielding a method capable of proving the general case. The next section develops this method.

3.8 Forcing

We will first recall the more general classes of formulas, $\forall_n$ and $\exists_n$, which we will use from this point forward. This will make our ultimate results more general.

In attempting to iteratively construct models of infinitary sentences using elementary chains as in the $n = 4$ and $n = 5$ cases, one is faced with the obstruction that infinitary formulas are not preserved by elementary extensions, meaning that work that has been done to ensure the truth of a subformula can be undone in the next extension. In fact, infinitary sentences can be very unstable with respect to elementary extensions.

Theorem 3.19. *There is a sentence $\psi \in \mathcal{L}_{\omega_1,\omega}$ and a structure $\mathcal{A}$ such that for any $\mathcal{B} \succeq \mathcal{A}$ with $\mathcal{B} \models \psi$, there is a $\mathcal{C} \succ \mathcal{B}$ with $\mathcal{C} \models \neg\psi$, and for any $\mathcal{B} \succeq \mathcal{A}$ with $\mathcal{B} \models \neg\psi$, there is a $\mathcal{C} \succ \mathcal{B}$ with $\mathcal{C} \models \psi$.*

Proof. Let $\mathcal{L}$ consist of a unary relation symbol Q, a binary relation symbol R, and for each natural number, a unary relation P_n. The structures we will be interested in will be the disjoint unions of certain building blocks. Think of elements satisfying Q as the roots of a block, and R attaching a number of other elements to the root. A "standard block" consists of a single element a satisfying Q, and countably many elements $b_0, b_1, b_2, \ldots$ not satisfying Q. We let $(a, b_0), (a, b_1), \dots \in R$. For each n, b_n satisfies P_n, but not P_m for $m \neq n$. A "non-standard block" consists of the same elements but, in addition, one or more "non-standard elements" b_* that do not satisfy P_n for any n.

Let $\mathcal{A}$ be the disjoint union of countably copies of the standard block. Any elementary extension $\mathcal{B}$ of $\mathcal{A}$ is elementarily equivalent to $\mathcal{A}$, and so consists of the disjoint union of infinitely many blocks which are either standard, or contain non-standard elements.

Let

$$\psi = \forall x \left(Q(x) \to \exists y \left(R(x, y) \wedge \bigwedge_n \neg P_n(y) \right) \right)$$

The sentence ψ says that every element satisfying Q has an associated non-standard element, so belongs to a non-standard block.

Let $\mathcal{B}$ be an elementary extension of $\mathcal{A}$. If $\mathcal{B} \models \psi$, we obtain $\mathcal{C} \succ \mathcal{B}$ by adding a single standard block, in which case $\mathcal{C} \models \neg\psi$. If $\mathcal{B} \models \neg\psi$, we obtain $\mathcal{C} \succ \mathcal{B}$ by adding a single non-standard element to each standard block, in which case $\mathcal{C} \models \psi$. $\square$

To solve the problem that this phenomenon raises for constructing models, we define relations between structures and infinitary sentences that keep track of our ability to make formulas true in further elementary extensions. These relations can be thought of as a notion of forcing, where in lieu of forcing conditions, we use structures, ordered by elementary extension. This is the approach taken by Robinson in [Rob71], and can be viewed as an extension of those methods to infinitary languages, where all extensions considered are elementary.

3.8.1 The Strong Forcing Relation

Given a structure $\mathcal{A}$, $\psi(\overline{x}) \in \mathcal{L}_{\infty,\omega}$, and $\overline{a} \in A$, we define the strong forcing relation $\mathcal{A} \Vdash \psi(\overline{a})$ by the following recursive clauses.

1. If ψ is atomic, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if $\mathcal{A} \models \psi(\overline{a})$.
2. If $\psi(\overline{x}) = \neg\phi(\overline{x})$, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for every $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \nVdash \phi(\overline{a})$.
3. If $\psi(\overline{x}) = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for some $\phi \in \Phi$, $\mathcal{A} \Vdash \phi(\overline{a})$.
4. If $\psi(\overline{x}) = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for every $\mathcal{B} \succ \mathcal{A}$, and $\phi \in \Phi$, there is a $\mathcal{C} \succ \mathcal{B}$ such that $\mathcal{C} \Vdash \phi(\overline{a})$.
5. If $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy})$, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for some $\overline{b} \in \mathcal{A}$, $\mathcal{A} \Vdash \phi(\overline{ab})$.
6. If $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for every $\mathcal{B} \succ \mathcal{A}$ and $\overline{b} \in \mathcal{B}$, there is a $\mathcal{C} \succ \mathcal{B}$ such that $\mathcal{C} \Vdash \phi(\overline{ab})$.

This diverges from the definition of the satisfaction relation in clauses (2), (4) and (6). For finitary formulas, this makes no difference.

Lemma 3.20. *If ψ is finitary, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if $\mathcal{A} \models \psi(\overline{a})$.*

Proof. We will prove this by induction on the complexity of ψ. All cases except those covered by clauses (2), (4) and (6) are identical to the satisfaction relation. For clause (2), let $\psi(\overline{x}) = \neg\phi(\overline{x})$. If ϕ is finitary, then $\mathcal{A} \Vdash \neg\phi(\overline{a})$ if and only if for every $\mathcal{B} \prec \mathcal{A}$, $\mathcal{B} \nVdash \phi(\overline{a})$. Appealing to induction, this is true if and only if for every $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \models \neg\phi(\overline{a})$, which is true if and only if $\mathcal{A} \models \neg\phi(\overline{a})$.

For clause (4), let $\psi = \bigwedge_{\phi\in\Phi} \phi(\overline{x})$, where Φ is finite. $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for each $\phi \in \Phi$ and $\mathcal{B} \succ \mathcal{A}$, there is a $\mathcal{C} \succ \mathcal{B}$ such that $\mathcal{C} \Vdash \phi(\overline{a})$, or appealing to induction, $\mathcal{C} \models \phi(\overline{a})$. Because ϕ is finitary, this is true if and only if for every such $\mathcal{B}$, $\mathcal{B} \models \phi(\overline{a})$, or equivalently, if $\mathcal{A} \models \phi(\overline{a})$, for each $\phi \in \Phi$. This, in turn, is true if and only if $\mathcal{A} \models \psi(\overline{a})$.

For clause (6), let $\psi = \forall\overline{y}\phi(\overline{xy})$. If ϕ is finitary, $\mathcal{A} \Vdash \forall\overline{y}\phi(\overline{ay})$ if and only if for every $\mathcal{B} \succ \mathcal{A}$, and every $\overline{b} \in \mathcal{B}$, there is a $\mathcal{C} \succ \mathcal{B}$ such that $\mathcal{C} \Vdash \phi(\overline{ab})$. Appealing to induction, this is true if and only if for every $\mathcal{B} \succ \mathcal{A}$, $\overline{b} \in \mathcal{B}$, there is a $\mathcal{C} \succ \mathcal{B}$ such that $\mathcal{C} \models \phi(\overline{ab})$. Because ϕ is finitary, this is true if and only if for every $\mathcal{B} \succ \mathcal{A}$, $\overline{b} \in \mathcal{B}$, $\mathcal{B} \models \phi(\overline{ab})$, or equivalently, for every $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \models \forall\overline{y}\phi(\overline{ay})$. This, in turn, is true if and only if $\mathcal{A} \models \forall\overline{y}\phi(\overline{ay})$. $\square$

For infinitary formulas, the relation $\Vdash$ is more stable than the satisfaction relation with respect to elementary extensions, as Theorem 3.19 and the following lemma demonstrate.

Lemma 3.21. *If $\mathcal{A} \prec \mathcal{B}$, and $\mathcal{A} \Vdash \psi(\overline{a})$, then $\mathcal{B} \Vdash \psi(\overline{a})$.*

Proof. We will prove this by induction on the complexity of ψ. If ψ is atomic, this is trivial. If $\psi = \neg\phi$, and $\mathcal{A} \Vdash \psi(\overline{a})$, then for any $\mathcal{C} \succ \mathcal{B}$, $\mathcal{C} \succ \mathcal{A}$, so $\mathcal{C} \nVdash \phi(\overline{a})$. Therefore, $\mathcal{B} \Vdash \psi(\overline{a})$. If $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$, this follows by induction. If $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$, and $\mathcal{A} \Vdash \psi(\overline{a})$, then for any $\mathcal{C} \succ \mathcal{B}$, and $\phi \in \Phi$, $\mathcal{C} \succ \mathcal{A}$, so there is a $\mathcal{D} \succ \mathcal{C}$ such that $\mathcal{D} \Vdash \phi(\overline{a})$. Therefore, $\mathcal{B} \Vdash \psi(\overline{a})$. If $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy})$, and $\mathcal{A} \Vdash \psi(\overline{a})$, then for some $\overline{b} \in \mathcal{A}$, $\mathcal{A} \Vdash \phi(\overline{ab})$. Appealing to induction, $\mathcal{B} \Vdash \phi(\overline{ab})$, so $\mathcal{B} \Vdash \psi(\overline{a})$. If $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$, and $\mathcal{A} \Vdash \psi(\overline{a})$. Suppose $\mathcal{C} \succ \mathcal{B}$, and $\overline{c} \in \mathcal{C}$. Then, $\mathcal{C} \succ \mathcal{A}$, so because $\mathcal{A} \Vdash \psi(\overline{a})$, there is a $\mathcal{D} \succ \mathcal{C}$ such that $\mathcal{D} \Vdash \phi(\overline{ac})$. Therefore, $\mathcal{B} \Vdash \psi(\overline{a})$. $\square$

Example 3.22. In Theorem 3.19, $\mathcal{A} \Vdash \psi$, so this is true for all elementary extensions of $\mathcal{A}$ as well.

Proof. Recall that

$$\psi = \forall x \left(Q(x) \to \exists y \left(R(x,y) \wedge \bigwedge\!\!\!\bigwedge_n \neg P_n(y) \right) \right)$$

Suppose $\mathcal{B} \succ \mathcal{A}$, and $a \in \mathcal{B}$. It suffices to show that for some $\mathcal{C} \succ \mathcal{A}$

$$\mathcal{C} \Vdash \neg Q(a) \vee \exists y (R(a,x) \wedge \bigwedge\!\!\!\bigwedge_n \neg P_n(y))$$

In the case that $\mathcal{B} \models \neg Q(a)$, we can take $\mathcal{C} = \mathcal{B}$, so it suffices to consider the case that $\mathcal{B} \models Q(a)$. In this case, let $\mathcal{C}$ be obtained by adding a non-standard element b, in the sense of Theorem 3.19 to the block corresponding to a. We will show that

$$\mathcal{C} \Vdash R(a,b) \wedge \bigwedge\!\!\!\bigwedge_n \neg P_n(b)$$

It suffices to show that $\mathcal{C} \Vdash R(a,b)$, and for each n, $\mathcal{C} \Vdash \neg P_n(b)$. This is true because $\mathcal{C} \models R(a,b)$, and $\mathcal{C} \models \neg P_n(b)$ for each n. □

For a structure $\mathcal{A}$, $\overline{a} \in \mathcal{A}$, and a formula ψ, it is immediate from the definition of $\mathcal{A} \Vdash \neg\phi(\overline{a})$ that it cannot be the case that $\mathcal{A} \Vdash \neg\phi(\overline{a})$ and $\mathcal{A} \Vdash \phi(\overline{a})$. We say that $\mathcal{A}$ decides $\phi(\overline{a})$ if either $\mathcal{A} \Vdash \psi(\overline{a})$ or $\mathcal{A} \Vdash \neg\psi(\overline{a})$.

Lemma 3.23. *For any structure $\mathcal{A}$, $\overline{a} \in \mathcal{A}$, and formula $\psi(\overline{x})$, there is a $\mathcal{B} \succ \mathcal{A}$ such that $\mathcal{B}$ decides $\psi(\overline{a})$.*

Proof. If $\mathcal{A} \Vdash \neg\psi(\overline{a})$, we can take $\mathcal{B} = \mathcal{A}$. Otherwise, there is some $\mathcal{B} \succ \mathcal{A}$ with $\mathcal{B} \Vdash \psi(\overline{a})$. □

3.8.2 Generic Structures

In order to obtain useful information from the forcing relation, we will construct structures in which formulas we have forced become true. This construction is largely independent of the formula under consideration, and the necessary property can be defined purely in terms of the forcing relation.

Let $\mathbb{A}$ be a fragment of $\mathcal{L}_{\infty,\omega}$. We say that a structure $\mathcal{G}$ is $\mathbb{A}$-generic if for any $\psi \in \mathbb{A}$, and $\overline{a} \in \mathcal{G}$, $\mathcal{G}$ decides $\psi(\overline{a})$. The next lemma shows that generic structures can be constructed, starting with any structure.

Lemma 3.24. *For any structure $\mathcal{A}$ and fragment $\mathbb{A}$, there is an $\mathbb{A}$-generic $\mathcal{G} \succ \mathcal{A}$.*

Proof. Let $\mathcal{C}$ be a structure. We will construct a structure $F(\mathcal{C})$ extending $\mathcal{C}$ as follows. Consider the set of pairs $(\psi(\overline{x}), \overline{c})$ with $\psi \in \mathbb{A}$, and $\overline{c} \in \mathcal{C}$ of length $\overline{x}$. Let $\{(\psi_\alpha(\overline{x}), \overline{c}_\alpha) | \alpha < \gamma\}$ be a well ordering of this set. We will define an elementary chain of length γ by transfinite recursion. Let $\mathcal{C}_0 = \mathcal{C}$. Having defined $\mathcal{C}_\alpha$, we define $\mathcal{C}_{\alpha+1}$ as follows. By Lemma 3.23, there is some $\mathcal{B} \succ \mathcal{C}_\alpha$ that decides $\psi_\alpha(\overline{c}_\alpha)$. Let $\mathcal{C}_{\alpha+1} = \mathcal{B}$. For limit ordinals $\beta < \gamma$, let $\mathcal{C}_\beta = \bigcup_{\alpha<\beta} \mathcal{C}_\alpha$. This defines an elementary chain $\{\mathcal{C}_\alpha | \alpha < \gamma\}$.

Let $F(\mathcal{C}) = \bigcup_{\alpha<\gamma} \mathcal{C}_\alpha$. Then, $F(\mathcal{C}) \succ \mathcal{C}$, and for every $\psi(\overline{x}) \in \mathbb{A}$, $\overline{c} \in \mathcal{C}$, $F(\mathcal{C})$ decides $\psi(\overline{c})$. Now consider the elementary chain

$$\mathcal{A} \prec F(\mathcal{A}) \prec F^2(\mathcal{A}) \prec \ldots$$

Let $\mathcal{G} = \bigcup_n F^n(\mathcal{A})$. Then, $\mathcal{A} \prec \mathcal{G}$. For any $\psi(\overline{x}) \in \mathbb{A}$, $\overline{a} \in \mathcal{G}$, $\overline{a} \in F^n(\mathcal{A})$ for some n, so $F^{n+1}(\mathcal{A})$ decides $\psi(\overline{a})$, which implies that $\mathcal{G}$ decides $\psi(\overline{a})$ because $F^{n+1}(\mathcal{A}) \prec \mathcal{G}$. Therefore, $\mathcal{G}$ is $\mathbb{A}$-generic. □

The next lemma shows that generic structures have the desired property, providing models of formulas we have forced.

Lemma 3.25. *Suppose $\psi(\overline{x}) \in \mathbb{A}$, $\mathcal{G}$ is $\mathbb{A}$-generic, and $\overline{a} \in \mathcal{G}$. Then, $\mathcal{G} \Vdash \psi(\overline{a})$ if and only if $\mathcal{G} \models \psi(\overline{a})$.*

Proof. We prove this by induction on the complexity of ψ. For ψ atomic, this is true by definition. Suppose $\psi(\overline{x}) = \neg\phi(\overline{x})$. Then $\mathcal{G} \Vdash \psi(\overline{a})$ if and only if $\mathcal{G} \nVdash \phi(\overline{a})$, because $\mathcal{G}$ is $\mathbb{A}$-generic. Appealing to induction, this is true if and only if $\mathcal{G} \not\models \phi(\overline{a})$, or equivalently, if $\mathcal{G} \models \neg\phi(\overline{a})$. For $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$, or $\psi = \exists\overline{y}\phi(\overline{y})$, the defining clause of $\Vdash$ is identical to that of the satisfaction relation, and the claim follows by induction.

Suppose $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$. If $\mathcal{G} \Vdash \psi(\overline{a})$, then for any ϕ, there is a $\mathcal{B} \succ \mathcal{G}$ such that $\mathcal{B} \Vdash \phi(\overline{a})$. Because $\mathcal{G}$ decides $\phi(\overline{a})$, it must be that $\mathcal{G} \Vdash \phi(\overline{a})$, so appealing to induction, $\mathcal{G} \models \phi(\overline{a})$. We conclude that $\mathcal{G} \models \psi(\overline{a})$. Suppose conversely that $\mathcal{G} \models \psi(\overline{a})$. Then, for each $\phi \in \Phi$, $\mathcal{G} \models \phi(\overline{a})$, so appealing to induction, $\mathcal{G} \Vdash \phi(\overline{a})$. For any $\mathcal{B} \succ \mathcal{G}$, Lemma 3.21 implies that $\mathcal{B} \Vdash \phi(\overline{a})$. We conclude that $\mathcal{G} \Vdash \psi(\overline{a})$.

Suppose $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$. If $\mathcal{G} \Vdash \psi(\overline{a})$, then for every $\overline{b} \in \mathcal{G}$, there is a $\mathcal{B} \succ \mathcal{G}$ such that $\mathcal{B} \Vdash \phi(\overline{ab})$. In this case, $\mathcal{G} \nVdash \neg\phi(\overline{ab})$, so $\mathcal{G} \Vdash \phi(\overline{ab})$. Appealing to induction, $\mathcal{G} \models \phi(\overline{ab})$. We conclude that $\mathcal{G} \models \psi(\overline{a})$. Suppose conversely that $\mathcal{G} \nVdash \psi(\overline{a})$. Then there is some $\mathcal{B} \succ \mathcal{G}$, and $\overline{b} \in \mathcal{B}$, such that for any $\mathcal{C} \succ \mathcal{B}$, $\mathcal{C} \nVdash \phi(\overline{ab})$. Let $\mathcal{G}' \succ \mathcal{B}$ be $\mathbb{A}$-generic. Then, $\mathcal{G}' \Vdash \neg\phi(\overline{ab})$. This implies that $\mathcal{G}' \Vdash \exists\overline{y}\neg\phi(\overline{ay})$, so $\mathcal{G} \Vdash \exists\overline{y}\neg\phi(\overline{ay})$. That is, there is some $\overline{b} \in \mathcal{G}$ such that $\mathcal{G} \Vdash \neg\phi(\overline{ab})$. In this case, $\mathcal{G} \nVdash \phi(\overline{ab})$, so appealing to induction, $\mathcal{G} \not\models \phi(\overline{ab})$. Then, $\mathcal{G} \not\models \psi(\overline{a})$. $\square$

It may seem as though we are making arbitrary choices about which formulas to force when we construct a generic extension of a structure $\mathcal{A}$. However, due to elementary amalgamation, these choices can only be made in one way.

Lemma 3.26. *Let $\mathcal{A}$ be a structure, $\overline{a} \in \mathcal{A}$, and $\psi(\overline{x})$ be a formula. If $\mathcal{A} \prec \mathcal{B}$, and $\mathcal{B} \Vdash \psi(\overline{a})$, then for every $\mathcal{C} \succ \mathcal{A}$ that decides $\psi(\overline{a})$, $\mathcal{C} \Vdash \psi(\overline{a})$.*

Proof. Suppose $\mathcal{A} \prec \mathcal{C}$ and $\mathcal{C} \Vdash \neg\psi(\overline{a})$. By the elementary amalgamation theorem, there is a $\mathcal{D}$ such that $\mathcal{B} \prec \mathcal{D}$, and $\mathcal{C} \prec \mathcal{D}$. Then, Lemma 3.21 implies that $\mathcal{D} \Vdash \psi(\overline{a})$, and $\mathcal{D} \Vdash \neg\psi(\overline{a})$, which is a contradiction. $\square$

Because of Lemma 3.26, we can regard all the information as to which formulas will be forced by extensions of a structure $\mathcal{A}$ as already present in $\mathcal{A}$. The next section defines a relation that captures this information.

3.8.3 The Weak Forcing Relation

It will be convenient to work with the weak forcing relation, denoted $\mathcal{A} \Vdash^* \psi(\overline{a})$. This can be defined in a variety of ways, all of which are equivalent. We will provisionally define $\mathcal{A} \Vdash^* \psi$ by the following recursive clauses.

1. If ψ is atomic, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \models \psi(\overline{a})$.
2. If $\psi(\overline{x}) = \neg\phi(\overline{x})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for every $\mathcal{B} \succ \mathcal{A}$, $\mathcal{A} \nVdash^* \phi(\overline{a})$.
3. If $\psi(\overline{x}) = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for some $\phi \in \Phi$, $\mathcal{A} \Vdash^* \phi(\overline{a})$.
4. If $\psi(\overline{x}) = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for every $\phi \in \Phi$, $\mathcal{A} \Vdash^* \phi(\overline{a})$.
5. If $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for some $\mathcal{B} \succ \mathcal{A}$ and $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash^* \phi(\overline{ab})$.
6. If $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for every $\mathcal{B} \succ \mathcal{A}$, and $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash^* \phi(\overline{ab})$.

The following lemma establishes other equivalent characterizations of the weak forcing relation.

Lemma 3.27. *Let ψ be an $\mathcal{L}_{\infty,\omega}$ formula. The following are equivalent.*

1. *$\mathcal{A} \Vdash^* \psi(\overline{a})$.*
2. *$\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$.*
3. *For some $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \Vdash \psi(\overline{a})$.*

4. *For every $\mathcal{B} \succ \mathcal{A}$ that decides $\psi(\overline{a})$, $\mathcal{B} \Vdash \psi(\overline{a})$*

5. *If $\psi \in \mathbb{A}$, and $\mathcal{G} \succ \mathcal{A}$ is $\mathbb{A}$-generic, $\mathcal{G} \models \psi(\overline{a})$.*

Proof. First, we will show that (2) implies (3). If $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$, then $\mathcal{A} \not\Vdash \neg\psi(\overline{a})$, so there is a $\mathcal{B} \succ \mathcal{A}$ such that $\mathcal{B} \Vdash \psi(\overline{a})$. That (3) implies (4) follows from Lemma 3.26.

Now, we will show that (4) implies (5). If $\psi \in \mathbb{A}$, and $\mathcal{G} \succ \mathcal{A}$ is $\mathbb{A}$-generic, then $\mathcal{G} \succ \mathcal{A}$ decides $\psi(\overline{a})$. If (4) holds, $\mathcal{G} \Vdash \psi(\overline{a})$, so by Lemma 3.25, $\mathcal{G} \models \psi(\overline{a})$.

Next, we will show that (5) implies (2). Suppose that if $\psi \in \mathbb{A}$ and $\mathcal{G} \succ \mathcal{A}$ is $\mathbb{A}$-generic, $\mathcal{G} \models \psi(\overline{a})$. By Lemma 3.25, $\mathcal{G} \Vdash \psi(\overline{a})$. Suppose $\mathcal{B} \succ \mathcal{A}$. By Lemma 3.23, there is some $\mathcal{C} \succ \mathcal{B}$ that decides $\psi(\overline{a})$. By Lemma 3.26, $\mathcal{C} \Vdash \psi(\overline{a})$. Consequently, for all $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \not\Vdash \neg\psi(\overline{a})$, so $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$.

Finally, to show that (1) and (2) are equivalent, we will show that the relation $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$ satisfies the recursive clauses of the definition of $\Vdash^*$. In doing so, we will use the equivalence of (2)-(5).

1. If ψ is atomic, then $\neg\neg\psi$ is finitary, so by Lemma 3.20 $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$ if and only if $\mathcal{A} \models \neg\neg\psi(\overline{a})$ if and only if $\mathcal{A} \models \psi(\overline{a})$.

2. If $\psi = \neg\phi$, $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$ if and only if, for every $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \not\Vdash \neg\psi(\overline{a})$. $\neg\psi = \neg\neg\phi$, so this is true if and only if for every $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \not\Vdash \neg\neg\phi(\overline{a})$.

3. Suppose $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$, and that $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$. Then, there is a $\mathcal{B} \succ \mathcal{A}$ such that $\mathcal{B} \Vdash \psi(\overline{a})$. Consequently, $\mathcal{A} \Vdash \neg\neg\phi(\overline{a})$. Suppose conversely that $A \Vdash \neg\neg\phi(\overline{a})$ for some $\phi \in \Phi$. Then, for some $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \Vdash \phi(\overline{a})$, so $\mathcal{B} \Vdash \psi(\overline{a})$. We conclude that $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$.

4. Suppose $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$. If $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$, then for some $\mathcal{B} \succ \mathcal{A}$, $\mathcal{B} \Vdash \psi(\overline{a})$. Then, for each $\phi \in \Phi$, there is a $\mathcal{C} \succ \mathcal{B}$ such that $\mathcal{C} \Vdash \phi(\overline{a})$. $\mathcal{C} \succ \mathcal{A}$ as well, so we conclude that $\mathcal{A} \Vdash \neg\neg\phi(\overline{a})$, for every $\phi \in \Phi$. Suppose conversely that for every $\phi \in \Phi$, $\mathcal{A} \Vdash \neg\neg\phi(\overline{a})$. Let $\mathbb{A}$ be a fragment containing ψ, and $\mathcal{G} \succ \mathcal{A}$ be $\mathbb{A}$-generic. Then, $\mathcal{G} \models \phi(\overline{a})$ for every $\phi \in \Phi$, so $\mathcal{G} \models \psi(\overline{a})$. We conclude that $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$.

5. Suppose $\psi = \exists\overline{y}\phi(\overline{y})$, and that $\mathcal{A} \Vdash \neg\neg\psi(\overline{x})$. Then, there is a $\mathcal{B} \succ \mathcal{A}$ such that $\mathcal{A} \Vdash \psi(\overline{a})$, so there is a $\overline{b} \in \mathcal{B}$ such that $\mathcal{B} \Vdash \phi(\overline{ab})$. $\mathcal{B} \succeq \mathcal{B}$, so we conclude that $\mathcal{B} \Vdash \neg\neg\phi(\overline{ab})$. Conversely, if for some $\mathcal{B} \succ \mathcal{A}$, $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash \neg\neg\phi(\overline{ab})$, then for some $\mathcal{C} \succ \mathcal{B}$, $\mathcal{C} \Vdash \phi(\overline{ab})$. Consequently, $\mathcal{C} \Vdash \psi(\overline{a})$. Therefore, $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$.

6. Suppose that $\psi = \forall\overline{y}\phi(\overline{y})$, and $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$. If $\mathcal{B} \succ \mathcal{A}$, and $\overline{b} \in \mathcal{B}$, then $\mathcal{B} \Vdash \neg\neg\psi(\overline{a})$, so there is some $\mathcal{C} \succ \mathcal{B}$ such that $\mathcal{C} \Vdash \psi(\overline{a})$. In this case, there is some $\mathcal{D} \succ \mathcal{C}$ such that $\mathcal{D} \Vdash \phi(\overline{ab})$. We conclude that $\mathcal{B} \Vdash \neg\neg\phi(\overline{ab})$. Suppose conversely that for any $\mathcal{B} \succ \mathcal{A}$, $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash \neg\neg\phi(\overline{a})$. Let $\mathbb{A}$ be a fragment containing ψ and $\mathcal{G} \succ \mathcal{A}$ be $\mathbb{A}$-generic. Then, for any $\overline{b} \in \mathcal{G}$, $\mathcal{G} \Vdash \neg\neg\phi(\overline{ab})$, so by Lemma 3.25, $\mathcal{G} \models \neg\neg\phi(\overline{ab})$. Consequently, $\mathcal{G} \models \phi(\overline{ab})$ for every $\overline{b} \in \mathcal{G}$, so $\mathcal{G} \models \psi(\overline{a})$. We conclude that $\mathcal{A} \models \neg\neg\psi(\overline{a})$. ☐

From the third characterization of the weak forcing relation, we have the following.

Corollary 3.28. *If $\mathcal{A} \Vdash \psi(\overline{a})$, then $\mathcal{A} \Vdash^* \psi(\overline{a})$.*

We can now establish some useful properties enjoyed by the weak forcing relation.

Lemma 3.29. *If $\mathcal{A} \prec \mathcal{B}$, then for any $\psi \in \mathcal{L}_{\infty,\omega}$, $\overline{a} \in \mathcal{A}$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{B} \Vdash^* \psi(\overline{a})$.*

Proof. Suppose that $\mathcal{A} \Vdash^* \psi(\overline{a})$. Then, if $\mathcal{C} \succ \mathcal{B}$, and $\mathcal{C}$ decides $\psi(\overline{a})$, $\mathcal{C} \succ \mathcal{A}$, so $\mathcal{C} \Vdash \psi(\overline{a})$. We conclude that $\mathcal{B} \Vdash^* \psi(\overline{a})$. Suppose now that $\mathcal{B} \Vdash^* \psi(\overline{a})$. Then, for some $\mathcal{C} \succ \mathcal{B}$, $\mathcal{C} \Vdash \psi(\overline{a})$. $\mathcal{C} \succ \mathcal{A}$, so $\mathcal{A} \Vdash^* \psi(\overline{a})$. ☐

Lemma 3.30. *For any $\mathcal{A}$, $\overline{a} \in \mathcal{A}$, and formula ψ, either $\mathcal{A} \Vdash^* \psi(\overline{a})$ or $\mathcal{A} \Vdash^* \neg\psi(\overline{a})$.*

Proof. By Lemma 3.23, there is some $\mathcal{B} \succ \mathcal{A}$ such that either $\mathcal{B} \Vdash \psi(\overline{a})$, or $\mathcal{B} \Vdash \neg\psi(\overline{a})$. In the first case, $\mathcal{A} \Vdash^* \psi(\overline{a})$, and in the second, $\mathcal{A} \Vdash^* \neg\psi(\overline{a})$ ☐

This implies that $\mathcal{A} \Vdash^* \neg\phi(\overline{a})$ if and only if $\mathcal{A} \not\Vdash^* \phi(\overline{a})$. Observing that recursive clauses (1) through (4) of the definition of the weak forcing relation are now identical to those of the satisfaction relation, we obtain the following.

Corollary 3.31. *If ψ is quantifier-free, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \models \psi(\overline{a})$.*

The following lemma shows that the weak forcing relation respects entailment and equivalence of formulas.

Lemma 3.32. *If $\psi_1 \vdash \psi_2$ and $\mathcal{A} \Vdash^* \psi_1(\overline{a})$, then $\mathcal{A} \Vdash^* \psi_2(\overline{a})$.*

Proof. Suppose that $\mathcal{A} \Vdash^* \psi_1(\overline{a})$ and $\mathcal{A} \not\Vdash^* \psi_2(\overline{a})$. By Lemma 3.30, $\mathcal{A} \Vdash^* \neg\psi_2(\overline{a})$. Let $\mathbb{A}$ be a fragment containing ψ_1 and $\neg\psi_2$, and let $\mathcal{G} \succ \mathcal{A}$ be $\mathbb{A}$-generic. Then, $\mathcal{G} \models \psi_1(\overline{a})$ and $\mathcal{G} \models \neg\psi_2(\overline{a})$, so $\psi_1 \not\vdash \psi_2$. $\square$

We can also define generic structures in terms of the weak forcing relation.

Lemma 3.33. *$\mathcal{G}$ is $\mathbb{A}$-generic if and only if for $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy}) \in \mathbb{A}$, $\overline{a} \in \mathcal{G}$, if $\mathcal{G} \Vdash^* \psi(\overline{a})$, there is a $\overline{b} \in \mathcal{G}$ such that $\mathcal{G} \Vdash^* \phi(\overline{ab})$.*

Proof. Suppose $\mathcal{G}$ is $\mathbb{A}$-generic. Let $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy}) \in \mathbb{A}$, and $\overline{a} \in \mathcal{G}$, If $\mathcal{G} \Vdash^* \psi(\overline{a})$, then $\mathcal{G} \models \psi(\overline{a})$, so for some $\overline{b} \in \mathcal{G}$, $\mathcal{G} \models \phi(\overline{ab})$. By Lemma 3.25, $\mathcal{G} \Vdash \phi(\overline{ab})$, so $\mathcal{G} \Vdash^* \phi(\overline{ab})$.

Conversely, suppose that for any $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy}) \in \mathbb{A}$, $\overline{a} \in \mathcal{G}$, if $\mathcal{G} \Vdash^* \psi(\overline{a})$, then for some $\overline{b} \in \mathcal{G}$, $\mathcal{G} \Vdash^* \phi(\overline{ab})$. We will show that $\mathcal{G}$ is $\mathbb{A}$-generic. By Lemma 3.30, it suffices to show that if $\psi \in \mathbb{A}$, and $\mathcal{G} \Vdash^* \psi(\overline{a})$, then $\mathcal{G} \Vdash \psi(\overline{a})$. We will prove this by induction on the complexity of ψ.

If ψ is atomic, then $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if $\mathcal{A} \models \psi(\overline{a})$. If $\psi = \neg\phi$, and $\mathcal{G} \Vdash^* \psi(\overline{a})$, then for every $\mathcal{B} \succ \mathcal{G}$, $\mathcal{B} \not\Vdash^* \phi(\overline{a})$. By Corollary 3.28, for every such $\mathcal{B}$, $\mathcal{B} \not\Vdash \phi(\overline{a})$, so $\mathcal{G} \Vdash \psi(\overline{a})$. If $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$, and $\mathcal{G} \Vdash^* \psi(\overline{a})$, then $\mathcal{G} \Vdash^* \phi(\overline{a})$ for some $\phi \in \Phi$. Appealing to induction, $\mathcal{G} \Vdash \phi(\overline{a})$, so $\mathcal{G} \Vdash \psi(\overline{a})$. Likewise, if $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$ and $\mathcal{G} \Vdash^* \psi(\overline{a})$, then $\mathcal{G} \Vdash^* \phi(\overline{a})$ for every $\phi \in \Phi$, so $\mathcal{G} \Vdash \phi(\overline{a})$ for every $\phi \in \Phi$, which implies that $\mathcal{G} \Vdash \psi(\overline{a})$.

Suppose that $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy})$. If $\mathcal{G} \Vdash^* \psi(\overline{a})$, then for some $\overline{b} \in \mathcal{G}$, $\mathcal{G} \Vdash^* \phi(\overline{ab})$. Appealing to induction, $\mathcal{G} \Vdash \phi(\overline{ab})$, so $\mathcal{G} \Vdash \psi(\overline{ab})$. Suppose $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$, and $\mathcal{G} \Vdash^* \psi(\overline{a})$. Then, for any $\mathcal{B} \succ \mathcal{G}$, $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash^* \phi(\overline{ab})$. Let $\mathcal{C} \succ \mathcal{B}$ decide $\phi(\overline{ab})$. Then, $\mathcal{C} \Vdash \phi(\overline{ab})$. We conclude that $\mathcal{G} \Vdash \psi(\overline{a})$. $\square$

The following lemma shows that the weak forcing relation depends only on first order properties.

Lemma 3.34. *Suppose $(\mathcal{A}, \overline{a}) \equiv (\mathcal{B}, \overline{b})$, then, for any formula ψ, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{B} \Vdash^* \psi(\overline{b})$.*

Proof. If $(\mathcal{A}, \overline{a}) \equiv (\mathcal{B}, \overline{a})$, the elementary amalgamation theorem implies that there is a structure $\mathcal{C}$ and elementary embeddings $f : \mathcal{A} \hookrightarrow \mathcal{C}$ and $g : \mathcal{B} \hookrightarrow \mathcal{C}$ such that $f(\overline{a}) = g(\overline{b})$. We can then identify $\mathcal{A}$ and $\mathcal{B}$ with elementary substructures of $\mathcal{C}$ so that $\overline{a} = \overline{b}$. Using Lemma 3.29, we have that $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{C} \Vdash^* \psi^*(\overline{a})$ if and only if $\mathcal{B} \Vdash^* \psi(\overline{b})$. $\square$

3.8.4 Definability

Recall that $\psi \in \mathcal{L}_{\infty,\omega}$ is elementary if it is of the form $\psi = \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta_{\alpha,\beta}$, for $\theta_{\alpha,\beta}$ finitary formulas. Lemma 3.34 shows that the weak forcing relation depends only on first order properties. The next lemma shows that is can be defined in terms of first order formulas.

Lemma 3.35. *For each $\psi \in \mathcal{L}_{\infty,\omega}$, there is an elementary formula* Force_ψ *such that $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \models \mathrm{Force}_\psi(\overline{a})$. Moreover, if ψ is a $\forall_n$ (resp. $\exists_n$) formula, then $\mathrm{Force}_\psi(\overline{a})$ can be taken to be a $\forall_n$ (resp. $\exists_n$) formula as well.*

Without the last clause, the lemma follows quite simply from Lemma 3.34. Consider the following set of types $\mathcal{T}$. Let $\mathcal{T} = \{\mathrm{tp}^{\mathcal{A}}(\bar{a}) : \mathcal{A} \Vdash^* \psi(\bar{a})\}$. By Lemma 3.34, $\mathcal{B} \Vdash^* \psi(\bar{a})$ if and only if, for some $p \in \mathcal{T}$, $\mathcal{B} \models p(\bar{a})$. Then let

$$\mathrm{Force}_\psi(x) = \bigvee\!\!\!\bigvee_{p(\bar{x})\in\mathcal{T}} \bigwedge\!\!\!\bigwedge_{\varphi\in p(\bar{x})} \varphi(\bar{x}).$$

However, we need a more involved argument if we want Force_ψ to have the same quantifier complexity as ψ.

Proof. We will define Force_ψ by recursion. At each step, we will ensure that Force_ψ is at most the complexity of ψ. If ψ is atomic, let $\text{Force}_\psi = \psi$. Suppose $\psi = \neg\phi$. By Lemma 3.30, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \nVdash^* \phi(\overline{a})$. Let $\text{Force}_\phi = \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta_{\alpha,\beta}$. Then, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \models \sim \text{Force}_\phi(\overline{a})$, the formal negation of Force_ϕ.

$$\sim \text{Force}_\phi = \bigwedge\!\!\!\bigwedge_\alpha \bigvee\!\!\!\bigvee_\beta \sim \theta_{\alpha,\beta}$$

which is equivalent to

$$\bigvee\!\!\!\bigvee_{f:\alpha\mapsto\beta} \bigwedge\!\!\!\bigwedge_\alpha \sim \theta_{\alpha,f(\alpha)}$$

We define Force_ψ to be this.

Suppose $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$. We can then define Force_ψ as $\bigvee\!\!\!\bigvee_{\phi\in\Phi} \text{Force}_\phi$. Suppose $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$. Then, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for every $\phi \in \Phi$, $\mathcal{A} \models \text{Force}_\phi(\overline{a})$. Let $\text{Force}_\phi = \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta^\phi_{\alpha,\beta}$. Then, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if

$$\mathcal{A} \models \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta^\phi_{\alpha,\beta}(\overline{a})$$

This formula is equivalent to

$$\bigvee\!\!\!\bigvee_{f:\phi\mapsto\alpha} \bigwedge\!\!\!\bigwedge_{\phi\in\Phi,\beta} \theta^\phi_{f(\phi),\beta}$$

We define Force_ψ to be this.

Suppose $\psi(\overline{x}) = \exists\overline{y}\phi(\overline{xy})$. Let $\text{Force}_\phi = \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta_{\alpha,\beta}$. The following are equivalent.

1. $\mathcal{A} \Vdash^* \psi(\overline{a})$;
2. For some $\mathcal{B} \succ \mathcal{A}$ and $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash^* \phi(\overline{ab})$;
3. For some $\mathcal{B} \succ \mathcal{A}$ and $\overline{b} \in \mathcal{B}$, $\mathcal{B} \models \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta_{\alpha,\beta}(\overline{ab})$;
4. For some $\mathcal{B} \succ \mathcal{A}$, $\overline{b} \in \mathcal{B}$, and α, $\mathcal{B} \models \theta_{\alpha,\beta}(\overline{ab})$ for each β;
5. For some α, the partial type $p_\alpha(\overline{y}) = \{\theta_{\alpha,\beta}(\overline{ay}) | \beta\}$ is finitely satisfiable in $\mathcal{A}$;
6. $$\mathcal{A} \models \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_{S\text{finite}} \exists\overline{y} \bigwedge_{\beta\in S} \theta_{\alpha,\beta}(\overline{ay}).$$

We define Force_ψ to be this formula. Suppose $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$. By Lemma 3.32, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \Vdash^* \neg\exists\overline{y}\neg\phi(\overline{ay})$, so we can use the rules for existential quantifiers and negations to construct $\text{Force}_\psi = \text{Force}_{\neg\exists\overline{y}\neg\phi(\overline{y})}$. ☐

A drawback of this definition is that for a cardinal κ, if $\psi \in \mathcal{L}_{\kappa,\omega}$, Force_ψ may not be in $\mathcal{L}_{\kappa,\omega}$. For instance, if $\mathcal{L}$ is countable, and $\psi \in \mathcal{L}_{\omega_1,\omega}$, Force_ψ may involve a disjunction over uncountably many formulas. The next results show that this drawback cannot be avoided.

Lemma 3.36. *There is a countable signature $\mathcal{L}$, and a sentence $\psi \in \mathcal{L}_{\omega_1,\omega}$ such that for any tree $T \subset \omega^{<\omega}$, there is a countable $\mathcal{L}$-structure $\mathcal{A}_T$, uniformly computable in T, satisfying $\mathcal{A}_T \Vdash^* \psi$ if and only if T has a path.*

Proof. Let $\mathcal{L}$ consist of unary relation symbols $R_{i,j}$ for $i, j \in \mathbb{N}$. Let $\psi = \exists x \bigwedge\!\!\!\bigwedge_i \bigvee\!\!\!\bigvee_j R_{i,j}(x)$. For a tree $T \subset \omega^{<\omega}$, we define $\mathcal{A}_T$ as follows. For each $\sigma \in T$, there is an element of $\mathcal{A}_T$ satisfying exactly the relations $R_{i,\sigma(i)}$ for each i less than the length of σ.

If $\mathcal{A}_T \Vdash^* \psi$, then for some $\mathcal{B} \succ \mathcal{A}_T$, and $b \in \mathcal{B}$, $\mathcal{B} \Vdash^* \bigwedge\limits_i \bigvee\limits_j R_{i,j}(b)$. In this case, $\mathcal{B} \models \bigwedge\limits_i \bigvee\limits_j R_{i,j}(b)$. Then, for some function $f \in \omega^\omega$, $\mathcal{B} \models R_{i,f(i)}(b)$ for each i. This implies that the partial type $\{R_{i,f(i)} | i < \omega\}$ is finitely satisfiable in $\mathcal{A}_T$, so for every n, there is a $a \in \mathcal{A}_T$ such that $\mathcal{A}_T \models R_{i,f(i)}(a)$ for $i < n$. That is, $f \restriction n \in T$, for all n, so f is a path in T. Suppose conversely that f is a path in T. Then, the partial type $\{R_{i,f(i)} | i < \omega\}$ is finitely satisfiable in $\mathcal{A}_T$, so for some elementary extension $\mathcal{B} \succ \mathcal{A}_T$, there is a $b \in \mathcal{B}$ realizing this type. Then, $\mathcal{B} \models \bigwedge\limits_i R_{i,f(i)}(B)$, so $\mathcal{B} \models \bigwedge\limits_i \bigvee\limits_j R_{i,j}(b)$. In this case, $\mathcal{B} \Vdash^* \bigwedge\limits_i \bigvee\limits_j R_{i,j}(b)$, so $\mathcal{A}_T \Vdash^* \psi$. □

Let $\mathrm{Mod}_{\mathcal{L}}$ be the set of ω-presentations of $\mathcal{L}$ structures. The mapping $T \mapsto \mathcal{A}_T$ witnesses the following.

Corollary 3.37. *The set $\{\mathcal{A} \in \mathrm{Mod}_{\mathcal{L}} | \mathcal{A} \Vdash^* \psi\}$ is $\mathbf{\Sigma}^1_1$ hard.*

We conclude that this set is not Borel, so is not the set of models of a $\mathcal{L}_{\omega_1,\omega}$ sentence. As such, we cannot have $\mathrm{Force}_\psi \in \mathcal{L}_{\omega_1,\omega}$.

Remark 3.38 (Structures of Bounded Cardinality). The apparatus built up in the previous sections can be adapted to consider only structures of cardinality below a particular bound κ. In the recursive definitions of the strong and weak forcing relations, one replaces elementary extensions in general with those of cardinality below κ. In order to construct generic structures of cardinality below κ, one also needs that the fragment $\mathbb{A}$ satisfies $|\mathbb{A}| < \kappa$, and so consists of $\mathcal{L}_{\kappa,\omega}$ formulas. Otherwise, the proofs go through without any changes. For instance, we can consider only countable structures, and countable fragments of $\mathcal{L}_{\omega_1,\omega}$, as is natural in applications to computable model theory.

3.9 The General Case

In this section, we use the forcing apparatus to prove the general case of Proposition 2.8.

Theorem 3.39. *Let T be a finitary theory. Let ψ be an infinitary ($\mathcal{L}_{\infty,\omega}$) $\forall_n$ formula which is equivalent to a finitary formula φ in all models of T. Then, ψ and φ are equivalent to a finitary $\forall_n$ formula in all models of T.*

To prove this, we will show that from the perspective of the weak forcing relation, $\forall_{n+1}$ formulas satisfy the semantic test of Theorem 3.15.

Lemma 3.40. *Suppose $\mathcal{A} \prec_n \mathcal{B}$ and $\overline{a} \in \mathcal{A}$. Let ψ be a $\forall_{n+1}$ formula. Then, either $\mathcal{A} \Vdash^* \psi(\overline{a})$ or $\mathcal{B} \Vdash^* \neg\psi(\overline{a})$.*

We will give two proofs of this fact. The first is syntactic, using the formula Force_ψ.

Proof. Let $\mathrm{Force}_\psi = \bigvee_\alpha \bigwedge_\beta \theta_{\alpha,\beta}$, where each $\theta_{\alpha,\beta}$ is a finitary $\forall_{n+1}$ formula. If $\mathcal{B} \nVdash^* \neg\psi(\overline{a})$, then by Lemma 3.30, $\mathcal{B} \Vdash^* \psi(\overline{a})$, so $\mathcal{B} \models \mathrm{Force}_\psi(\overline{a})$. Then, for some α, and every β, $\mathcal{B} \models \theta_{\alpha,\beta}(\overline{a})$. Because each $\theta_{\alpha,\beta}$ is a finitary $\forall_{n+1}$ formula, $\mathcal{A} \models \theta_{\alpha,\beta}(\overline{a})$ for the same α, and every β. Therefore, $\mathcal{A} \models \mathrm{Force}_\psi(\overline{a})$, so $\mathcal{A} \Vdash^* \psi(\overline{a})$. We conclude that either $\mathcal{A} \Vdash^* \psi(\overline{a})$, or $\mathcal{B} \Vdash^* \neg\psi(\overline{a})$. □

The second proof uses the recursive definition of the weak forcing relation.

Proof. We will prove this by induction on n, and on the complexity of ψ. For $n = 0$, we have $\mathcal{A} \subset \mathcal{B}$. Because ψ is $\forall_1$, it has one of the following forms.

1. $\psi = \neg\phi$, where ϕ is $\exists_1$.
2. $\psi = \bigvee_{\phi \in \Phi} \phi$, where each $\phi \in \Phi$ is $\forall_1$.
3. $\psi = \bigwedge_{\phi \in \Phi} \phi$, where each $\phi \in \Phi$ is $\forall_1$.
4. $\psi = \forall \overline{y} \phi(\overline{y})$, for $\phi(\overline{y})$ a $\forall_1$ formula.

5. $\psi = \forall \overline{y} \phi(\overline{y})$, for $\phi(\overline{y})$ quantifier free.

In case (1), Lemma 3.32 allows us to replace ψ with $\sim \phi$, the formal negation of ϕ, which is less complex than ψ, so we can appeal to induction. In cases (2) and (3), we can similarly appeal to induction. In case (4), suppose $\mathcal{B} \nVdash^* \neg\psi(\overline{a})$, so that $\mathcal{B} \Vdash \psi(\overline{a})$. Suppose $\mathcal{A}' \succ \mathcal{A}$, and $\overline{b} \in \mathcal{A}'$. By Lemma 3.16, there is a $\mathcal{B}' \succ \mathcal{B}$ such that $\mathcal{A}' \subset \mathcal{B}'$. Because $\mathcal{B} \Vdash^* \psi(\overline{a})$, $\mathcal{B}' \Vdash^* \phi(\overline{ab})$. Appealing to induction, $\mathcal{A}' \Vdash^* \phi(\overline{ab})$. We conclude that $\mathcal{A} \Vdash^* \psi(\overline{a})$. Therefore, either $\mathcal{B} \Vdash^* \neg\psi(\overline{a})$ or $\mathcal{A} \Vdash^* \psi(\overline{a})$. Case (5) can be handled similarly. Rather than appealing to induction, we have that because ϕ is quantifier free, $\mathcal{B}' \Vdash^* \phi(\overline{ab})$ if and only if $\mathcal{B}' \models \phi(\overline{ab})$. As shown in Corollary 3.4, this is true if and only if $\mathcal{A}' \models \phi(\overline{ab})$, which is true if and only if $\mathcal{A}' \Vdash^* \phi(\overline{ab})$. Note that by taking $\overline{y}$ to be variables not present in ϕ, and applying Lemma 3.32, case (5) covers quantifier free formulas as well.

The inductive step is very similar. Suppose ψ is $\forall_{n+1}$, for $n > 0$. As before, ψ has one of the following forms.

1. $\psi = \neg\phi$, where ϕ is $\exists_{n+1}$.
2. $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$, where each $\phi \in \Phi$ is $\forall_{n+1}$.
3. $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$, where each $\phi \in \Phi$ is $\forall_{n+1}$.
4. $\psi = \forall \overline{y} \phi(\overline{y})$, for $\phi(\overline{y})$ a $\forall_{n+1}$ formula.
5. $\psi = \forall \overline{y} \phi(\overline{y})$, for $\phi(\overline{y})$ a $\exists_n$ formula.

Cases (1) through (4) can be handled by appealing to induction on the complexity of ψ as before. In case (5), suppose that $\mathcal{A}' \succ \mathcal{A}$, and $\overline{b} \in \mathcal{A}'$. By Lemma 3.16, there is a $\mathcal{B}' \succ \mathcal{B}$ such that $\mathcal{A}' \prec_n \mathcal{B}'$. By Lemma 3.17, there is a $\mathcal{C}$ such that $\mathcal{B}' \prec_{n-1} \mathcal{C}$ and $\mathcal{A}' \prec \mathcal{C}$. Appealing to induction on n, either $\mathcal{B}' \Vdash^* \sim \phi(\overline{ab})$, or $\mathcal{C} \Vdash^* \neg \sim \phi(\overline{ab})$. In the first case, $\mathcal{B} \Vdash^* \sim \psi(\overline{a})$, so by Lemma 3.32, $\mathcal{B} \Vdash^* \neg\psi(\overline{a})$. In the second case, Lemma 3.32 implies that $\mathcal{C} \Vdash^* \phi(\overline{ab})$, so Lemma 3.29 implies that $\mathcal{A}' \Vdash^* \phi(\overline{ab})$. If this happens for every $\mathcal{A}'$, $\overline{b}$, then $\mathcal{A} \Vdash^* \psi(\overline{a})$. Otherwise, $\mathcal{B} \Vdash^* \neg\psi(\overline{a})$.

□

With Lemma 3.40, we can now prove Theorem 3.39.

Proof. It suffices to consider $n \geq 1$, as $n = 0$ is covered by Corollary 3.4. We will prove the contrapositive. Suppose that φ is not equivalent to any finitary $\forall_n$ formula over T. As shown in Theorem 3.15, this implies that there are models of T, $\mathcal{A} \prec_{n-1} \mathcal{B}$, and $\overline{a} \in \mathcal{A}$, such that $\mathcal{A} \models \neg\varphi(\overline{a})$ and $\mathcal{B} \models \varphi(\overline{a})$. By Lemma 3.40, either $\mathcal{A} \Vdash^* \psi(\overline{a})$ or $\mathcal{B} \Vdash^* \neg\psi(\overline{a})$. Let $\mathbb{A}$ be a fragment containing ψ. In the first case, let $\mathcal{G}$ be $\mathbb{A}$-generic, such that $\mathcal{A} \prec \mathcal{G}$. In the second, let $\mathcal{H}$ be $\mathbb{A}$-generic, such that $\mathcal{B} \prec \mathcal{H}$. In the first case, $\mathcal{G} \models T + \neg\varphi(\overline{a}) + \psi(\overline{a})$, and in the second, $\mathcal{H} \models T + \varphi(\overline{a}) + \neg\psi(\overline{a})$. We conclude that φ and ψ are not equivalent over T.

□

Remark 3.41. In addition to the results on relative decidability, this also can be used to connect the complexity of finitary sentences to descriptive set theoretic notions of complexity. Suppose that $\mathcal{L}$ is a countable and φ is a sentence of $\mathcal{L}_{\omega,\omega}$. By the Löwenheim-Skolem theorem for $\mathcal{L}_{\omega_1,\omega}$, if ψ is a sentence of $\mathcal{L}_{\omega_1,\omega}$ and φ and ψ are equivalent in all countable structures, they are equivalent in all structures. Using Theorem 3.39 and Vaught's version of the Lopez-Escobar theorem [Vau75], we have that the following are equivalent.

1. φ is equivalent to a finitary $\forall_n$ sentence (respectively $\exists_n$).
2. $\{\mathcal{A} \in \mathrm{Mod}_{\mathcal{L}} | \mathcal{A} \models \varphi\}$ is $\mathbf{\Pi}^0_n$ (respectively, $\mathbf{\Sigma}^0_n$).

References

[AKMS89] C. Ash, J. Knight, M. Manasse, and T. Slaman. Generic copies of countable structures. *Ann. Pure Appl. Logic*, 1989.

[Bau06] A. Bauer. König's lemma and the Kleene tree, 2006.

[Cha59] C. C. Chang. On unions of chains of models. *Proc. Amer. Math. Soc.*, 10:120–127, 1959.

[Chi90] J. Chisholm. Effective model theory vs. recursive model theory. *Journal of Symbol Logic*, 1990.

[CMS21] J. Chubb, R. Miller, and R. Solomon. Model completeness and relative decidability. *Archive for Mathematical Logic*, 2021.

[ŁS55] J. Łoś and R. Suszko. On the infinite sums of models. *Bull. Acad. Polon. Sci. Cl. III.*, 3:201–202, 1955.

[Mar16] David Marker. *Lectures on Infinitary Model Theory*. Lecture Notes in Logic. Cambridge University Press, 2016.

[Mon21] Antonio Montalbán. *Computable Structure Theory: Beyond the arithmetic*. Unpublished Draft, 2021.

[Rob71] Abraham Robinson. Infinite forcing in model theory. In *Proceedings of the Second Scandinavian Logic Symposium (Oslo, 1970)*, pages 317–340. Studies in Logic and the Foundations of Mathematics, Vol. 63, 1971.

[Vau75] Robert Vaught. Invariant sets in topology and logic. *Fund. Math.*, 82:269–294, 1974/75.