Infinitary Logic Has No Expressive Efficiency Over Finitary Logic

Matthew Harrison-Trainor Miles Kretschmer

October 23, 2022

Abstract

We can measure the complexity of a logical formula by counting the number of alternations between existential and universal quantifiers. Suppose that an elementary first-order formula φ (in $\mathcal{L}_{\omega,\omega}$) is equivalent to a formula of the infinitary language $\mathcal{L}_{\infty,\omega}$ with n alternations of quantifiers. We prove that φ is equivalent to a finitary formula with n alternations of quantifiers. Thus using infinitary logic does not allow us to express a finitary formula in a simpler way.

1 Introduction

This paper is about the relationship between finitary elementary first-order logic ($\mathcal{L}_{\omega,\omega}$) and the infinitary logics (such as $\mathcal{L}_{\infty,\omega}$) which extend it by allowing conjunctions and disjunctions of infinite sets of formulas. These infinitary logics are more expressive than the finitary logic, but lose compactness. For example, there are classes such as connected graphs and torsion groups which cannot be axiomatized in the finitary logic (as shown by a simple compactness argument) but which can be axiomatized in the infinitary logic.

A natural way of measuring the complexity of a formula is by putting it in normal form and counting the number of alternations of quantifiers. An $\exists_n$ formula is a formula which begins with a block of existential quantifiers, and has n alternating blocks of existential and universal quantifiers. Similarly, a $\forall_n$ formula has n alternating blocks of existential and universal quantifiers, beginning with a block of universal quantifiers. In the infinitary languages, we count alternations of quantifiers, but we do *not* count infinitary conjunctions and disjunctions. (This differs from the standard way of counting quantifiers in computable structure theory, where formulas are classified as Σ_n or Π_n; a Σ_n formula is $\exists_n$ but not necessarily vice versa.[1])

Now suppose that we have a finitary formula $\varphi(\bar{x})$ which is equivalent to an infinitary existential ($\exists_1$) formula, say

$$\psi(\bar{x}) = \bigwedge_i \bigvee_j \exists \bar{y}_{i,j}\, \theta_{i,j}(\bar{x}, \bar{y}_{i,j}).$$

[1] Our theorems remain true if one replaces $\exists_n$ by Σ_n and $\forall_n$ by Π_n, so we get stronger theorems by not counting infinitary conjunctions and disjunctions.

Now one can check that $\psi(\bar{x})$ (and hence $\varphi(\bar{x})$) has the property of being preserved in extensions: if $\mathcal{A} \subseteq \mathcal{B}$, and $\mathcal{A} \models \psi(\bar{a})$, then $\mathcal{B} \models \psi(\bar{a})$. Since $\varphi(\bar{x})$ is preserved by extensions, a standard preservation result says that $\varphi(\bar{x})$ is itself equivalent to a finitary existential ($\exists_1$) formula. In general, one proves that if a finitary formula is equivalent to an infinitary existential formula, it is equivalent to a finitary existential formula.

If $\varphi(\bar{x})$ is instead a finitary formula which is equivalent to an infinitary $\forall_2$ formula $\psi(\bar{x})$, a similar argument works to show that $\varphi(\bar{x})$ is equivalent to a finitary $\forall_2$ formula. Instead of preservation in extensions, we use the fact that $\varphi(\bar{x})$ is equivalent to a finitary $\forall_2$ formula if and only φ is preserved under unions of chains of models: Whenever $\mathcal{A}_0 \subseteq \mathcal{A}_1 \subseteq \mathcal{A}_2 \subseteq \cdots \subseteq \mathcal{A}_\omega = \bigcup_{n<\omega} \mathcal{A}_n$ is a chain of models, with $\mathcal{A}_n \models \varphi(\bar{a})$ for all n, then $\mathcal{A}_\omega \models \psi(\bar{a})$. In this case, $\varphi(\bar{x})$ is preserved under unions of chains of models because it is equivalent to the infinitary $\forall_2$ formula $\psi(\bar{x})$, which one can check is preserved under unions of chains of models.

This sort of argument breaks down at the $\forall_4/\exists_4$ case, as there is no longer a semantic test for finitary $\forall_4/\exists_4$ formulas which is satisfied by infinitary $\forall_4/\exists_4$ formulas. Nevertheless, the main result of this paper is that this result is true in general.

Theorem 1.1. *Let T be a finitary ($\mathcal{L}_{\omega,\omega}$) theory. Let ψ be an infinitary ($\mathcal{L}_{\infty,\omega}$) $\exists_n$ (resp. $\forall_n$) formula which is equivalent to a finitary formula φ in all models of T. Then, ψ and φ are equivalent to a finitary $\exists_n$ (resp. $\forall_n$) formula in all models of T.*

In spite of their greater expressive power, infinitary languages cannot define relations already definable in finitary first order logic with any greater efficiency (as measured by quantifier complexity). One can also view this as computing an intersection: the properties which are expressible by both a finitary formula and an infinitary $\exists_n$ formula are exactly the properties expressible by a finitary $\exists_n$ formula. (One might compare the spirit of this result to Louveau's theorem [Lou80]: If $C \subset \omega^\omega$ is hyperarithmetic and $\mathbf{\Sigma}^0_\alpha$, for $\alpha < \omega_1^{\mathrm{CK}}$, then C is Σ^0_α.)

The general technique we use is a notion of forcing where the conditions are elementary extensions of a given structure. This can also be viewed as a form of Robinson's model-theoretic forcing [Rob71]. One can perhaps view the forcing as fixing a deficiency of infinitary formulas; while the truth of an infinitary formula φ may not be preserved by elementary extensions, forcing will be: if $\mathcal{A} \Vdash^* \varphi$ and $\mathcal{B} \succeq \mathcal{A}$, then we will have that $\mathcal{B} \Vdash^* \varphi$. If φ is equivalent to a finitary formula, then the truth of φ is already preserved under elementary extensions, and we will be able to show that $\mathcal{A} \models \varphi$ if and only if $\mathcal{A} \Vdash^* \varphi$.

The definability of the forcing relation will allow us to characterize those infinitary formulas which do transfer across elementary extensions. We will show that these are the infinitary formulas built up using conjunctions and disjunctions of finitary formulas.

Theorem 1.2. *Let $\psi(\bar{x})$ be an infinitary $\mathcal{L}_{\infty,\omega}$ formula. The following are equivalent:*

(1) Given $\mathcal{A} \preceq \mathcal{B}$, $\mathcal{A} \models \varphi(\bar{a})$ if and only if $\mathcal{B} \models \varphi(\bar{a})$.

(2) $\psi(\bar{x})$ is equivalent to an $\mathcal{L}_{\infty,\omega}$ formula of the form

$$\bigvee_{\alpha} \bigwedge_{\beta} \theta_{\alpha,\beta}(\bar{x})$$

where each $\theta_{\alpha,\beta}$ is a finitary formula.

Moreover, if these conditions hold and ψ is a $\forall_n$ (resp. $\exists_n$) formula, then we may take each $\theta_{\alpha,\beta}$ to be $\forall_n$ (resp. $\exists_n$).

We call a formula as in (2) an *elementary formula.* One particular consequence is that if an infinitary formula φ is preserved under elementary extensions, and is $\exists_{n+1}$, then whenever $\mathcal{A} \models \varphi$ and $\mathcal{A} \leq_n \mathcal{B}$, $\mathcal{B} \models \varphi$.

Note that even if $\psi(\bar{x})$ is in $\mathcal{L}_{\omega_1,\omega}$, the formula $\bigvee_\alpha \bigwedge_\beta \theta_{\alpha,\beta}(\bar{x})$ might not be. However, we show that one can find a formula in $\mathcal{L}_{\omega_1,\omega}$ witnessing that $\psi(\bar{x})$ is preserved upwards and downwards by elementary extensions. In the following theorem, we say that an infinitary formula is *quantifier-free over finitary formulas* if it can be built by repeatedly taking (infinitary) conjunctions, disjunctions, and negations of finitary formulas.

Theorem 1.3. *Let $\psi(\bar{x})$ be an infinitary $\mathcal{L}_{\omega_1,\omega}$ formula. The follow are equivalent:*

(1) Given $\mathcal{A} \preceq \mathcal{B}$, $\mathcal{A} \models \psi(\bar{a})$ if and only if $\mathcal{B} \models \psi(\bar{a})$.

(2) $\psi(\overline{x})$ is equivalent to an $\mathcal{L}_{\omega_1,\omega}$ formula which is quantifier-free over finitary formulas.

Moreover, if these conditions hold and ψ is a $\forall_n$ formula (or an $\exists_n$ formula), then $\psi(\overline{x})$ is equivalent to an $\mathcal{L}_{\omega_1,\omega}$ formula which is quantifier-free over finitary $\exists_n/\forall_n$ formulas.

We leave open whether this generalizes to $\mathcal{L}_{\kappa,\omega}$ for $\kappa > \omega_1$.

2 Preliminaries

2.1 The Infinitary Languages $\mathcal{L}_{\kappa,\omega}$ and $\mathcal{L}_{\infty,\omega}$

Fix a signature τ and κ an infinite cardinal. The language $\mathcal{L}_{\kappa,\omega}(\tau)$, which we will habitually denote $\mathcal{L}_{\kappa,\omega}$, will be the language which allows infinite conjunctions and disjunctions of size $< \kappa$. We define the language $\mathcal{L}_{\kappa,\omega}$ to be the smallest set of formulas with the following properties.

(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 $\phi \in \mathcal{L}_{\kappa,\omega}$, then $\forall x\phi \in \mathcal{L}_{\kappa,\omega}$ and $\exists x\phi \in \mathcal{L}_{\kappa,\omega}$.

(4) If $\Phi \subset \mathcal{L}_{\kappa,\omega}$ share finitely many free variables and $|\Phi| < \kappa$, then $\bigvee\limits_{\phi\in\Phi} \phi \in \mathcal{L}_{\kappa,\omega}$ and $\bigwedge\limits_{\phi\in\Phi} \phi \in \mathcal{L}_{\kappa,\omega}$

Note that $\mathcal{L}_{\omega,\omega}$ is just the standard finitary elementary first-order logic. We say that a formula is in $\mathcal{L}_{\infty,\omega}$ if it is a formula of $\mathcal{L}_{\kappa,\omega}$ for some κ. We often refer to the formulas of $\mathcal{L}_{\omega,\omega}$ as finitary formulas and to those of $\mathcal{L}_{\infty,\omega}$ as infinitary formulas.

A fragment $\mathbb{A}$ of $\mathcal{L}_{\kappa,\omega}$ is a set of formulas of $\mathcal{L}_{\kappa,\omega}$ 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 $\mathbb{A}$.

Starting with any 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.

2.2 $\forall_\alpha$ and $\exists_\alpha$ Formulas

In order to count quantifier alternations in infinitary formulas, we will define a hierarchy of classes of formulas, ranked by natural numbers. For each $n \in \mathbb{N}$, we define the classes $\forall_n$, $\exists_n$ of formulas of $\mathcal{L}_{\infty,\omega}$ as follows.

(1) If ψ is atomic, then for all n, $\psi \in \forall_n$, and $\psi \in \exists_n$.

(2) If $\psi = \neg\phi$, then $\psi \in \exists_n$, (respectively, $\psi \in \forall_n$) if $\phi \in \forall_n$ (respectively, $\phi \in \exists_n$).

(3) If $\psi = \bigvee\!\!\!\bigvee_{\phi\in\Phi} \phi$ or $\psi = \bigwedge\!\!\!\bigwedge_{\phi\in\Phi} \phi$, then $\psi \in \exists_n$ (respectively, $\psi \in \forall_n$) if for every $\phi \in \Phi$, $\phi \in \exists_n$ (respectively, $\phi \in \forall_n$).

(4) If $\psi = \exists\overline{y}\phi(\overline{y})$, $\psi \in \exists_n$ if $\phi \in \exists_n$ and $n \geq 1$, and $\psi \in \exists_{n+1}$ if $\phi \in \forall_n$.

(5) If $\psi = \forall\overline{y}\phi(\overline{y})$, $\psi \in \forall_n$ if $\phi \in \forall_n$ and $n \geq 1$, and $\psi \in \forall_{n+1}$ if $\phi \in \exists_n$.

For example, $\exists_0 = \forall_0$ consists of quantifier free formulas. We call the $\exists_1$ formulas *existential formulas* and the $\forall_1$ formulas *universal formulas.*

We note that this way of counting quantifiers differs from the standard way of counting quantifiers for $\mathcal{L}_{\omega_1,\omega}$ formulas in computable structure theory. That is, an $\exists_n$ formula is not necessarily Σ_n and a $\forall_n$ formula is not necessarily Π_n (though every Σ_n formula is $\exists_n$ and every Π_n formula is $\forall_n$). The difference is that for Σ_n and Π_n formulas, we count infinite disjunctions the same as existential quantifiers, and infinite conjunctions the same as universal quantifiers, whereas for $\exists_n$ and $\forall_n$ formulas we do not count infinite conjunctions and disjunctions at all. So, for example, a formula of the form $\bigwedge\!\!\!\bigwedge_i \exists x\theta_i(x)$ is $\exists_1$ but not Σ_1.

There are natural reasons to consider both forms of counting. The classes Σ_n and Π_n have descriptive-set-theoretic meaning via Vaught's version of the Lopez-Escobar theorem [Vau75, LE65]: An invariant set $\mathbb{K}$ of structures is $\mathbf{\Sigma}^0_\alpha$ if and only if it is defined by a Σ_α formula. Moreover, many connections between definability and computability-theoretic properties work with Σ_α and Π_α formulas.

On the other hand, Malitz [Mal69] showed that a formula of $\mathcal{L}_{\omega_1,\omega}$ is preserved by substructures if and only if it is universal ($\forall_1$). This is the sort of behaviour we will consider in this paper, and so the classes $\forall_n$ and $\exists_n$ are the appropriate ones to use. We note that because every Π_n formula is also $\forall_n$ (and there is no difference for finitary formulas), our results about $\forall_n$ formulas apply to Π_n formulas as well. Stating our results for $\forall_n$ formulas is simply their strongest form.

3 Forcing with Elementary Extensions

In attempting to iteratively construct models of infinitary sentences using elementary chains, one is faced with the obstruction that infinitary formulas are not preserved by elementary extensions. In fact, infinitary sentences can be very unstable with respect to elementary extensions.

Theorem 3.1. *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 τ consist of a unary relation symbol Q, a binary relation symbol R, and a unary relation P_n for each natural number 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 $b_0, b_1, b_2, \dots$ 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

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

The sentence ψ says that every element satisfying Q has an associated non-standard element, and 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 $\mathcal{C} \models \psi$. □

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 as forcing conditions we use structures, ordered by elementary extension. This is related to 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.

As usual, there is a strong forcing and a weak forcing. The strong forcing is useful for defining genericity, while the weak notion is required for the definability of forcing.

3.1 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 $\mathcal{B} \succeq \mathcal{A}$, $\mathcal{B} \nVdash \phi(\overline{a})$.

(3) If $\psi(\overline{x}) = \bigvee_{\phi \in \Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for some $\phi \in \Phi$, $\mathcal{A} \Vdash \phi(\overline{a})$.

(4) If $\psi(\overline{x}) = \bigwedge_{\phi \in \Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if for every $\mathcal{B} \succeq \mathcal{A}$, and $\phi \in \Phi$, there is a $\mathcal{C} \succeq \mathcal{B}$ 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 $\mathcal{B} \geq \mathcal{A}$ and $\overline{b} \in \mathcal{B}$, there is a $\mathcal{C} \geq \mathcal{B}$ 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 3.2. *If ψ is finitary, $\mathcal{A} \Vdash \psi(\overline{a})$ if and only if $\mathcal{A} \vDash \psi(\overline{a})$.*

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 $\mathcal{B} \leq \mathcal{A}$, $\mathcal{B} \nVdash \phi(\overline{a})$. Appealing to induction, this is true if and only if for every $\mathcal{B} \geq \mathcal{A}$, $\mathcal{B} \vDash \neg\phi(\overline{a})$, which is true if and only if $\mathcal{A} \vDash \neg\phi(\overline{a})$.

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 $\mathcal{B} \geq \mathcal{A}$, there is a $\mathcal{C} \geq \mathcal{B}$ such that $\mathcal{C} \Vdash \phi(\overline{a})$, or appealing to induction, $\mathcal{C} \vDash \phi(\overline{a})$. Because ϕ is finitary, this is true if and only if for every such $\mathcal{B}$, $\mathcal{B} \vDash \phi(\overline{a})$, or equivalently, if $\mathcal{A} \vDash \phi(\overline{a})$, for each $\phi \in \Phi$. This, in turn, is true if and only if $\mathcal{A} \vDash \psi(\overline{a})$.

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 $\mathcal{B} \geq \mathcal{A}$, and every $\overline{b} \in \mathcal{B}$, there is a $\mathcal{C} \geq \mathcal{B}$ such that $\mathcal{C} \Vdash \phi(\overline{ab})$. Appealing to induction, this is true if and only if for every $\mathcal{B} \geq \mathcal{A}$, $\overline{b} \in \mathcal{B}$, there is a $\mathcal{C} \geq \mathcal{B}$ such that $\mathcal{C} \vDash \phi(\overline{ab})$. Because ϕ is finitary, this is true if and only if for every $\mathcal{B} \geq \mathcal{A}$, $\overline{b} \in \mathcal{B}$, $\mathcal{B} \vDash \phi(\overline{ab})$, or equivalently, for every $\mathcal{B} \geq \mathcal{A}$, $\mathcal{B} \vDash \forall\overline{y}\phi(\overline{ay})$. This, in turn, is true if and only if $\mathcal{A} \vDash \forall\overline{y}\phi(\overline{ay})$. □

For infinitary formulas, the relation $\Vdash$ is more stable than the satisfaction relation with respect to elementary extensions. (Contrast the Lemma below with Theorem 3.1.)

Lemma 3.3. *If $\mathcal{A} \leq \mathcal{B}$, 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 $\mathcal{C} \geq \mathcal{B}$, $\mathcal{C} \geq \mathcal{A}$, 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 $\mathcal{C} \geq \mathcal{B}$, and $\phi \in \Phi$, $\mathcal{C} \geq \mathcal{A}$, so there is a $\mathcal{D} \geq \mathcal{C}$ 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 $\mathcal{C} \geq \mathcal{B}$, and $\overline{c} \in \mathcal{C}$. Then, $\mathcal{C} \geq \mathcal{A}$, so because $\mathcal{A} \Vdash \psi(\overline{a})$, there is a $\mathcal{D} \geq \mathcal{C}$ such that $\mathcal{D} \Vdash \phi(\overline{ac})$. Therefore, $\mathcal{B} \Vdash \psi(\overline{a})$. □

Example 3.4. Let $\mathcal{A}$ and ψ be the structure and formula respectively from Theorem 3.1. Then $\mathcal{A} \Vdash \psi$ (and 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 $\mathcal{B} \geq \mathcal{A}$, and $a \in \mathcal{B}$. It suffices to show that for some $\mathcal{C} \geq \mathcal{A}$

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

In the case that $\mathcal{B} \vDash \neg Q(a)$, we can take $\mathcal{C} = \mathcal{B}$, so it suffices to consider the case that $\mathcal{B} \vDash Q(a)$. In this case, let $\mathcal{C}$ be obtained by adding a non-standard element b, in the sense of Theorem 3.1 to the block corresponding to a. We will show that

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

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 $\mathcal{C} \vDash R(a,b)$, and $\mathcal{C} \vDash \neg P_n(b)$ for each n. □

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

Lemma 3.5. *For any structure $\mathcal{A}$, $\overline{a} \in \mathcal{A}$, and formula $\psi(\overline{x})$, there is a $\mathcal{B} \geq \mathcal{A}$ 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 $\mathcal{B} \geq \mathcal{A}$ with $\mathcal{B} \Vdash \psi(\overline{a})$. □

3.1.1 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 3.6. *For any structure $\mathcal{A}$ and fragment $\mathbb{A}$, there is an $\mathbb{A}$-generic $\mathcal{G} \geq \mathcal{A}$.*

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 3.5, there is some $\mathcal{B} \geq \mathcal{C}_\alpha$ 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, $F(\mathcal{C}) \geq \mathcal{C}$, 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

$$\mathcal{A} \leq F(\mathcal{A}) \leq F^2(\mathcal{A}) \leq \dots$$

Let $\mathcal{G} = \bigcup_n F^n(\mathcal{A})$. Then, $\mathcal{A} \leq \mathcal{G}$. 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 $F^{n+1}(\mathcal{A}) \leq \mathcal{G}$. 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 3.7. *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 $\mathcal{G} \vDash \psi(\overline{a})$.*

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 $\mathcal{G} \nvDash \phi(\overline{a})$, or equivalently, if $\mathcal{G} \vDash \neg\phi(\overline{a})$. 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 $\mathcal{B} \geq \mathcal{G}$ 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, $\mathcal{G} \vDash \phi(\overline{a})$. We conclude that $\mathcal{G} \vDash \psi(\overline{a})$. Suppose conversely that $\mathcal{G} \vDash \psi(\overline{a})$. Then, for each $\phi \in \Phi$, $\mathcal{G} \vDash \phi(\overline{a})$, so appealing to induction, $\mathcal{G} \Vdash \phi(\overline{a})$. For any $\mathcal{B} \geq \mathcal{G}$, Lemma 3.3 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 $\mathcal{B} \geq \mathcal{G}$ 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, $\mathcal{G} \vDash \phi(\overline{ab})$. We conclude that $\mathcal{G} \vDash \psi(\overline{a})$. Suppose conversely that $\mathcal{G} \nVdash \psi(\overline{a})$. Then there is some $\mathcal{B} \geq \mathcal{G}$, and $\overline{b} \in \mathcal{B}$, such that for any $\mathcal{C} \geq \mathcal{B}$, $\mathcal{C} \nVdash \phi(\overline{ab})$. Let $\mathcal{G}' \geq \mathcal{B}$ 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, $\mathcal{G} \nvDash \phi(\overline{ab})$. Then, $\mathcal{G} \nvDash \psi(\overline{a})$. □

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. We highlight this as one of the key special features of this forcing.

Lemma 3.8. *Let $\mathcal{A}$ be a structure, $\overline{a} \in \mathcal{A}$, and $\psi(\overline{x})$ be a formula. If $\mathcal{B} \geq \mathcal{A}$, and $\mathcal{B} \Vdash \psi(\overline{a})$, then for every $\mathcal{C} \geq \mathcal{A}$ that decides $\psi(\overline{a})$, $\mathcal{C} \Vdash \psi(\overline{a})$.*

Proof. Suppose $\mathcal{A} \leq \mathcal{C}$ and $\mathcal{C} \Vdash \neg\psi(\overline{a})$. By the elementary amalgamation theorem, there is a $\mathcal{D}$ such that $\mathcal{B} \leq \mathcal{D}$, and $\mathcal{C} \leq \mathcal{D}$. Then, Lemma 3.3 implies that $\mathcal{D} \Vdash \psi(\overline{a})$, and $\mathcal{D} \Vdash \neg\psi(\overline{a})$, which is a contradiction. □

Because of Lemma 3.8, 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.

3.2 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 $\mathcal{A} \Vdash^* \psi$ to be $\mathcal{A} \Vdash \neg\neg\psi$. It is a standard fact that the weak forcing relation satisfies the following recursive clauses:

(1) If ψ is atomic, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \vDash \psi(\overline{a})$.

(2) If $\psi(\overline{x}) = \neg\phi(\overline{x})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for every $\mathcal{B} \geq \mathcal{A}$, $\mathcal{B} \nVdash^* \phi(\overline{a})$.

(3) If $\psi(\overline{x}) = \bigvee_{\phi\in\Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for each $\mathcal{B} \geq \mathcal{A}$, there is $\mathcal{C} \geq \mathcal{B}$ and $\phi \in \Phi$ such that $\mathcal{C} \Vdash^* \phi(\overline{a})$.

(4) If $\psi(\overline{x}) = \bigwedge_{\phi\in\Phi} \phi(\overline{x})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ 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 all $\mathcal{B} \geq \mathcal{A}$ there is $\mathcal{C} \geq \mathcal{B}$ and $\overline{b} \in \mathcal{C}$ such that $\mathcal{C} \Vdash^* \phi(\overline{ab})$.

(6) If $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for every $\overline{b} \in \mathcal{A}$, $\mathcal{A} \Vdash^* \phi(\overline{ab})$.

The following lemma establishes other equivalent characterizations of the weak forcing relation, which are unique to this notion of forcing and result from elementary amalgamation. (In essence, elementary amalgamation allows us to simplify being "dense below".)

Lemma 3.9. *Let ψ be an $\mathcal{L}_{\infty,\omega}$ formula. The following are equivalent.*

(1) $\mathcal{A} \Vdash^* \psi(\overline{a})$.

(2) For some $\mathcal{B} \geq \mathcal{A}$, $\mathcal{B} \Vdash \psi(\overline{a})$.

(3) For every $\mathcal{B} \geq \mathcal{A}$ that decides $\psi(\overline{a})$, $\mathcal{B} \Vdash \psi(\overline{a})$

(4) If $\psi \in \mathbb{A}$, and $\mathcal{G} \geq \mathcal{A}$ is $\mathbb{A}$-generic, $\mathcal{G} \vDash \psi(\overline{a})$.

Proof. First, we will show that (1) implies (2). If $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$, then $\mathcal{A} \nVdash \neg\psi(\overline{a})$, so there is a $\mathcal{B} \geq \mathcal{A}$ such that $\mathcal{B} \Vdash \psi(\overline{a})$. That (2) implies (3) follows from Lemma 3.8.

Now, we will show that (3) implies (4). If $\psi \in \mathbb{A}$, and $\mathcal{G} \geq \mathcal{A}$ is $\mathbb{A}$-generic, then $\mathcal{G} \geq \mathcal{A}$ decides $\psi(\overline{a})$. If (3) holds, $\mathcal{G} \Vdash \psi(\overline{a})$, so by Lemma 3.7, $\mathcal{G} \vDash \psi(\overline{a})$.

Next, we will show that (4) implies (1). Suppose that if $\psi \in \mathbb{A}$ and $\mathcal{G} \geq \mathcal{A}$ is $\mathbb{A}$-generic, $\mathcal{G} \vDash \psi(\overline{a})$. By Lemma 3.7, $\mathcal{G} \Vdash \psi(\overline{a})$. Suppose $\mathcal{B} \geq \mathcal{A}$. By Lemma 3.5, there is some $\mathcal{C} \geq \mathcal{B}$ that decides $\psi(\overline{a})$. By Lemma 3.8, $\mathcal{C} \Vdash \psi(\overline{a})$. Consequently, for all $\mathcal{B} \geq \mathcal{A}$, $\mathcal{B} \nVdash \neg\psi(\overline{a})$, so $\mathcal{A} \Vdash \neg\neg\psi(\overline{a})$. □

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

Corollary 3.10. *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 3.11. *If $\mathcal{A} \leq \mathcal{B}$, then for any $\psi \in \mathcal{L}_{\infty,\omega}$, and $\overline{a} \in \mathcal{A}$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{B} \Vdash^* \psi(\overline{a})$.*

Proof. Suppose that $\mathcal{A} \Vdash^* \psi(\overline{a})$. Then, if $\mathcal{C} \geq \mathcal{B}$, and $\mathcal{C}$ decides $\psi(\overline{a})$, $\mathcal{C} \geq \mathcal{A}$, 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 $\mathcal{C} \geq \mathcal{B}$, $\mathcal{C} \Vdash \psi(\overline{a})$. $\mathcal{C} \geq \mathcal{A}$, so $\mathcal{A} \Vdash^* \psi(\overline{a})$. □

As a consequence of elementary amalgamation, we get the following lemma which is an important and special property of this forcing notion.

Lemma 3.12. *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 3.5, there is some $\mathcal{B} \geq \mathcal{A}$ 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 $\mathcal{A} \nVdash^* \phi(\overline{a})$. Moreover, using Lemma 3.11 and elementary amalgamation:

Theorem 3.13. *The weak forcing $\Vdash^*$ is defined by the following recursive conditions:*

(1) If ψ is atomic, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \vDash \psi(\overline{a})$.

(2′) If $\psi(\overline{x}) = \neg\phi(\overline{x})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only $\mathcal{A} \nVdash^* \phi(\overline{a})$.

(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 there is $\mathcal{B} \geq \mathcal{A}$ and $\overline{b} \in \mathcal{B}$ such that $\mathcal{B} \Vdash^* \phi(\overline{ab})$.

(6) If $\psi(\overline{x}) = \forall\overline{y}\phi(\overline{xy})$, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if for every $\overline{b} \in \mathcal{A}$, $\mathcal{A} \Vdash^* \phi(\overline{ab})$.

Proof. As noted previously, $\Vdash^*$ satisfies clauses (1), (4) and (6). That $\Vdash^*$ satisfies (2′) follows immediately from Lemma 3.12.

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

For (5′), suppose $\psi = \exists\overline{y}\phi(\overline{y})$, and that $\mathcal{A} \Vdash^* \psi(\overline{x})$. Then, there is a $\mathcal{B} \geq \mathcal{A}$ such that $\mathcal{A} \Vdash \psi(\overline{a})$, so there is a $\overline{b} \in \mathcal{B}$ such that $\mathcal{B} \Vdash \phi(\overline{ab})$, in which case $\mathcal{B} \Vdash^* \phi(\overline{ab})$. Conversely, if for some $\mathcal{B} \geq \mathcal{A}$, $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash^* \phi(\overline{ab})$, then for some $\mathcal{C} \geq \mathcal{B}$, $\mathcal{C} \Vdash \phi(\overline{ab})$. In this case, $\mathcal{C} \Vdash \psi(\overline{a})$ so $\mathcal{A} \Vdash^* \psi(\overline{a})$. ☐

Note that (1), (2′), (3′), (4), and (6) are the same as the conditions for the satisfaction relation; it is only (5′) that differs. Only the quantifiers, and not the infinitary connectives, are treated differently; this is why we use the $\exists_n/\forall_n$ hierarchy rather than the Σ_n/Π_n hierarchy. Moreover, we obtain the following.

Corollary 3.14. *If ψ is quantifier-free, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \vDash \psi(\overline{a})$.*

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

Lemma 3.15. *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 $\mathcal{A} \nVdash^* \psi_2(\overline{a})$. By Lemma 3.12, $\mathcal{A} \Vdash^* \neg\psi_2(\overline{a})$. Let $\mathbb{A}$ be a fragment containing ψ_1 and $\neg\psi_2$, and let $\mathcal{G} \geq \mathcal{A}$ be $\mathbb{A}$-generic. Then, $\mathcal{G} \vDash \psi_1(\overline{a})$ and $\mathcal{G} \vDash \neg\psi_2(\overline{a})$, so $\psi_1 \nvDash \psi_2$. ☐

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

Lemma 3.16. *$\mathcal{G}$ 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 $\mathcal{G} \vDash \psi(\overline{a})$, so for some $\overline{b} \in \mathcal{G}$, $\mathcal{G} \vDash \phi(\overline{ab})$. By Lemma 3.7, $\mathcal{G} \Vdash \phi(\overline{ab})$, 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 3.12, 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 $\mathcal{A} \vDash \psi(\overline{a})$. If $\psi = \neg\phi$, and $\mathcal{G} \Vdash^* \psi(\overline{a})$, then for every $\mathcal{B} \geq \mathcal{G}$, $\mathcal{B} \nVdash^* \phi(\overline{a})$. By Corollary 3.10, for every such $\mathcal{B}$, $\mathcal{B} \nVdash \phi(\overline{a})$, 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 $\mathcal{B} \geq \mathcal{G}$, $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash^* \phi(\overline{ab})$. Let $\mathcal{C} \geq \mathcal{B}$ decide $\phi(\overline{ab})$. Then, $\mathcal{C} \Vdash \phi(\overline{ab})$. We conclude that $\mathcal{G} \Vdash \psi(\overline{a})$. □

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

Lemma 3.17. *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 3.11, we have that $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{C} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{B} \Vdash^* \psi(\overline{b})$. □

3.2.1 Definability

Recall that we said that a formula $\psi \in \mathcal{L}_{\infty,\omega}$ 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 3.17 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 3.18. *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 $\mathcal{A} \vDash \mathrm{Force}_\psi(\overline{a})$. Moreover, if ψ is a $\forall_n$ (resp. $\exists_n$) formula, then* $\mathrm{Force}_\psi(\overline{a})$ *can be taken to be of the form*

$$\bigvee\!\!\!\bigvee_{\alpha} \bigwedge\!\!\!\bigwedge_{\beta} \theta_{\alpha,\beta}(\bar{x})$$

where each $\theta_{\alpha,\beta}$ is a finitary $\forall_n$ (resp. $\exists_n$) formula.

Without the last clause, the lemma follows quite simply from Lemma 3.17. Consider the following set of types $\mathcal{T}$. Let $\mathcal{T} = \{\text{tp}^{\mathcal{A}}(\bar{a}) : \mathcal{A} \Vdash^* \psi(\bar{a})\}$. By Lemma 3.17, $\mathcal{B} \Vdash^* \psi(\bar{a})$ if and only if, for some $p \in \mathcal{T}$, $\mathcal{B} \vDash p(\bar{a})$. Then let

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

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 (recalling Theorem 3.13 which gives simplified conditions for the weak forcing). 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 3.12, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \not\Vdash^* \phi(\overline{a})$. Let $\text{Force}_\phi = \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta_{\alpha,\beta}$. Then, $\mathcal{A} \Vdash^* \psi(\overline{a})$ if and only if $\mathcal{A} \vDash \sim \text{Force}_\phi(\overline{a})$, the formal negation of Force_ϕ. This is

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

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 $\phi \in \Phi$, $\mathcal{A} \vDash \text{Force}_\phi(\overline{a})$. 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

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

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 $\mathcal{B} \geq \mathcal{A}$ and $\overline{b} \in \mathcal{B}$, $\mathcal{B} \Vdash^* \phi(\overline{ab})$;

(3) For some $\mathcal{B} \geq \mathcal{A}$ and $\overline{b} \in \mathcal{B}$, $\mathcal{B} \vDash \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta_{\alpha,\beta}(\overline{ab})$;

(4) For some $\mathcal{B} \geq \mathcal{A}$, $\overline{b} \in \mathcal{B}$, and α, $\mathcal{B} \vDash \theta_{\alpha,\beta}(\overline{ab})$ 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)
$$\mathcal{A} \vDash \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_{S \text{ finite}} \exists\overline{y} \bigwedge_{\beta\in S} \theta_{\alpha,\beta}(\overline{ay}).$$

We define Force_ψ to be this formula. Suppose $\psi(\overline{x}) = \forall \overline{y}\phi(\overline{xy})$. By Lemma 3.15, $\mathcal{A} \Vdash^* \psi(\overline{a})$ 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 $\mathrm{Force}_\psi = \mathrm{Force}_{\neg\exists\overline{y}\neg\phi(\overline{y})}$. □

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 τ 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 cannot be avoided.

Lemma 3.19. *There is a countable signature τ, and a sentence $\psi \in \mathcal{L}_{\omega_1,\omega}(\tau)$ such that for any tree $T \subset \omega^{<\omega}$, there is a countable τ-structure $\mathcal{A}_T$, uniformly computable in T, satisfying $\mathcal{A}_T \Vdash^* \psi$ if and only if T has a path.*

Proof. Let τ 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 $\mathcal{B} \succeq \mathcal{A}_T$, and $b \in \mathcal{B}$, $\mathcal{B} \Vdash^* \bigwedge\!\!\!\bigwedge_i \bigvee\!\!\!\bigvee_j R_{i,j}(b)$. In this case, $\mathcal{B} \vDash \bigwedge\!\!\!\bigwedge_i \bigvee\!\!\!\bigvee_j R_{i,j}(b)$. Then, for some function $f \in \omega^\omega$, $\mathcal{B} \vDash R_{i,f(i)}(b)$ for each i. This implies that the partial type $\left\{R_{i,f(i)} \middle| i < \omega\right\}$ is finitely satisfiable in $\mathcal{A}_T$, so for every n, there is a $a \in \mathcal{A}_T$ such that $\mathcal{A}_T \vDash R_{i,f(i)}(a)$ 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 $\left\{R_{i,f(i)} \middle| i < \omega\right\}$ is finitely satisfiable in $\mathcal{A}_T$, so for some elementary extension $\mathcal{B} \succeq \mathcal{A}_T$, there is a $b \in \mathcal{B}$ realizing this type. Then, $\mathcal{B} \vDash \bigwedge\!\!\!\bigwedge_i R_{i,f(i)}(B)$, so $\mathcal{B} \vDash \bigwedge\!\!\!\bigwedge_i \bigvee\!\!\!\bigvee_j R_{i,j}(b)$. In this case, $\mathcal{B} \Vdash^* \bigwedge\!\!\!\bigwedge_i \bigvee\!\!\!\bigvee_j R_{i,j}(b)$, so $\mathcal{A}_T \Vdash^* \psi$. □

Let Mod_τ be the Polish space of ω-presentations of τ-structures. The mapping $T \mapsto \mathcal{A}_T$ witnesses the following.

Corollary 3.20. *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 $\mathrm{Force}_\psi \in \mathcal{L}_{\omega_1,\omega}$.

3.3 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.

4 Applications of the Forcing Notion

We now apply the forcing with elementary extension introduced in the previous section to prove the main theorems of this paper.

4.1 The Main Theorem

In this section, we will prove Theorem 1.1.

Theorem 1.1. *Let T be a finitary ($\mathcal{L}_{\omega,\omega}$) theory. Let ψ be an infinitary ($\mathcal{L}_{\infty,\omega}$) $\exists_n$ (resp. $\forall_n$) formula which is equivalent to a finitary formula φ in all models of T. Then, ψ and φ are equivalent to a finitary $\exists_n$ (resp. $\forall_n$) formula in all models of T.*

To prove this, we will use the fact that a finitary formula φ is equivalent to a $\exists_n$ formula over a theory T if and only if whenever $\mathcal{A} \preceq_{n-1} \mathcal{B}$ are models of T, $\overline{a} \in \mathcal{A}$, and $\mathcal{A} \vDash \varphi(\overline{a})$, then $\mathcal{B} \vDash \varphi(\overline{a})$. This generalises the fact that a finitary formula is equivalent to an existential formula if and only if it is preserved upwards under superstructures.

However, while an $\exists_1$ infinitary formula is preserved upwards under superstructures, it is not generally true that an $\exists_n$ formula is preserved upwards under $(n-1)$-elementary superstructures. Instead, we will show that they are preserved upwards under $(n-1)$-elementary superstructures if we consider weak forcing rather than satisfaction.

Lemma 4.1. *Suppose $\mathcal{A} \preceq_{n-1} \mathcal{B}$ and $\overline{a} \in \mathcal{A}$. Let ψ be an infinitary $\exists_n$ formula. Then if $\mathcal{A} \Vdash^* \psi(\overline{a})$ then $\mathcal{B} \Vdash^* \psi(\overline{a})$.*

Proof. Let $\text{Force}_\psi = \bigvee\!\!\!\bigvee_\alpha \bigwedge\!\!\!\bigwedge_\beta \theta_{\alpha,\beta}$, where each $\theta_{\alpha,\beta}$ is a finitary $\exists_n$ formula. Suppose that $\mathcal{A} \Vdash^* \psi(\overline{a})$. Then $\mathcal{A} \vDash \text{Force}_\psi(\overline{a})$, and so for some α, and every β, $\mathcal{A} \vDash \theta_{\alpha,\beta}(\overline{a})$. Because each $\theta_{\alpha,\beta}$ is a finitary $\exists_n$ formula, $\mathcal{B} \vDash \theta_{\alpha,\beta}(\overline{a})$ for the same α, and every β. Therefore, $\mathcal{B} \vDash \text{Force}_\psi(\overline{a})$, so $\mathcal{B} \Vdash^* \psi(\overline{a})$. □

With Lemma 4.1, we can now prove Theorem 1.1.

Proof of Theorem 1.1. We prove the $\exists_n$ case; the $\forall_n$ case can be obtained by taking negations.

Suppose that, as in the hypotheses of Theorem 1.1, ψ is an infinitary $\exists_n$ formula which is equivalent to a finitary formula φ in all models of T. We want to show that ψ and φ are equivalent to a finitary $\exists_n$ formula in all models of T. To do this, suppose that $\mathcal{A} \preceq_{n-1} \mathcal{B}$ are models of T, $\bar{a} \in A$, and $\mathcal{A} \vDash \varphi(\bar{a})$; we must show that $\mathcal{B} \vDash \varphi(\bar{a})$.

Now since $\mathcal{A} \vDash \varphi(\bar{a})$ and φ is finitary, by Lemma 3.2 we have that $\mathcal{A} \Vdash^* \varphi(\bar{a})$. Since φ and ψ are equivalent in $\mathcal{A}$, and forcing respects this (Lemma 3.15), we have that $\mathcal{A} \Vdash^* \psi(\bar{a})$. But we just proved in Lemma 4.1 that forcing an $\exists_n$ formula is preserved upwards under $n-1$-elementary superstructures, and so $\mathcal{B} \Vdash^* \psi(\bar{a})$. Using the same equivalences as before, we get that $\mathcal{B} \Vdash^* \varphi(\bar{a})$ and then that $\mathcal{B} \vDash \varphi(\bar{a})$. This completes the argument. □

Remark 4.2. Suppose that τ 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 1.1 and Vaught's version of the Lopez-Escobar theorem [Vau75], we have that the following are equivalent.

(1) φ is equivalent to a finitary $\exists_n$ sentence (respectively $\forall_n$).

(2) $\{\mathcal{A} \in \text{Mod}_\tau | \mathcal{A} \vDash \varphi\}$ is $\mathbf{\Sigma}^0_n$ (respectively, $\mathbf{\Pi}^0_n$).

Thus Vaught's version of the Lopez-Escobar theorem specialises to the case of finitary formulas.

The proof of Theorem 1.1 we gave above makes use of standard ideas that show up in forcing, like the definability of forcing. One can also give a more hands-on proof which has a different sort of explanatory power. The outline of this proof is as follows. We prove the contrapositive: Supposing that $\varphi \in \mathcal{L}_{\omega,\omega}$ is not equivalent to any finitary $\exists_n$ formula over T, we aim to produce a model witnessing that φ is not equivalent over T to some particular $\exists_n$ formula ψ. Using the fact that φ is not equivalent to any finitary $\exists_n$ formula, we can construct models of T, $\mathcal{A} \preceq_{n-1} \mathcal{B}$, such that for some $\overline{a} \in \mathcal{A}$, $\mathcal{A} \vDash \varphi(\overline{a})$ and $\mathcal{B} \vDash \neg\varphi(\overline{a})$. It suffices then to construct either an elementary extension of $\mathcal{A}$ modeling $\neg\psi(\overline{a})$ or an elementary extension of $\mathcal{B}$ modeling $\psi(\overline{a})$. Considering $\mathbb{A}$-generic elementary extensions of $\mathcal{A}$ and $\mathcal{B}$ for a fragment $\mathbb{A}$ containing ψ, it suffices to show that either $\mathcal{A} \Vdash^* \neg\psi(\overline{a})$ or that $\mathcal{B} \Vdash^* \psi(\overline{a})$. This is the content of Lemma 4.1. This lemma can be proved by a more semantic route, using the following amalgamation lemmas, both of which are applications of compactness.

Lemma 4.3. *Suppose $\mathcal{A} \preceq_n \mathcal{B}$ and $\mathcal{A} \preceq \mathcal{A}'$. There is a $\mathcal{B}'$ such that $\mathcal{B}' \succeq \mathcal{B}$ and $\mathcal{B}' \succeq_n \mathcal{A}'$.*

Lemma 4.4. *If $\mathcal{A} \preceq_{n+1} \mathcal{B}$, then there is a $\mathcal{C} \succeq_n \mathcal{B}$ with $\mathcal{C} \succeq \mathcal{A}$.*

We then proceed by induction on n. Appealing to induction on the complexity of ψ, we can reduce to the case that $\psi(\overline{x})$ is of the form $\exists\overline{y}\eta(\overline{xy})$, where η is $\forall_{n-1}$. Suppose $\mathcal{A} \Vdash^* \psi(\overline{a})$. Then, for some $\mathcal{A}' \succeq \mathcal{A}$, and $\overline{a'} \in \mathcal{A}'$, $\mathcal{A}' \Vdash \eta(\overline{aa'})$. Applying Lemma 4.3, we have a $\mathcal{B}' \succeq \mathcal{B}$, such that $\mathcal{A}' \preceq_{n-1} \mathcal{B}'$. Applying Lemma 4.4, there is a $\mathcal{C} \succeq_{n-2} \mathcal{B}'$ such that $\mathcal{C} \succeq \mathcal{A}'$. Then, $\mathcal{C} \Vdash^* \eta(\overline{aa'})$. Appealing to induction on n, in the case of the pair of structures $\mathcal{B}' \preceq_{n-2} \mathcal{C}$ and the $\exists_{n-1}$ formula $\neg\eta$, we have that $\mathcal{B}' \Vdash^* \eta(\overline{aa'})$, so $\mathcal{B} \Vdash^* \psi(\overline{a})$. Unraveling this argument, the recursive definition of the weak forcing relation guides a construction of a sequence of elementary extensions containing witnesses for subformulas of either ψ or $\neg\psi$.

4.2 Preservation by Elementary Extensions

4.2.1 Preservation of formulas in $\mathcal{L}_{\infty,\omega}$

We will now prove Theorem 1.2.

Theorem 1.2. *Let $\psi(\bar{x})$ be an infinitary $\mathcal{L}_{\infty,\omega}$ formula. The following are equivalent:*

(1) Given $\mathcal{A} \preceq \mathcal{B}$, $\mathcal{A} \vDash \varphi(\bar{a})$ if and only if $\mathcal{B} \vDash \varphi(\bar{a})$.

(2) $\psi(\overline{x})$ is equivalent to an $\mathcal{L}_{\infty,\omega}$ formula of the form

$$\bigvee_{\alpha} \bigwedge_{\beta} \theta_{\alpha,\beta}(\bar{x})$$

where each $\theta_{\alpha,\beta}$ is a finitary formula.

Moreover, if these conditions hold and ψ is a $\forall_n$ (resp. $\exists_n$) formula, then we may take each $\theta_{\alpha,\beta}$ to be $\forall_n$ (resp. $\exists_n$).

Proof. Suppose that ψ is equivalent to an elementary formula $\phi = \bigvee_{\alpha} \bigwedge_{\beta} \theta_{\alpha,\beta}$. It suffices to show that ϕ transfers across elementary extensions. Let $\mathcal{A} \leq \mathcal{B}$, $\overline{a} \in \mathcal{A}$. If $\mathcal{A} \vDash \phi(\overline{a})$, then for some α, and every β, $\mathcal{A} \vDash \theta_{\alpha,\beta}(\overline{a})$. Then, for the same α, and every β, $\mathcal{B} \vDash \theta_{\alpha,\beta}(\overline{a})$, so $\mathcal{B} \vDash \phi(\overline{a})$. Conversely, if $\mathcal{B} \vDash \phi(\overline{a})$, then for some α, and every β, $\mathcal{B} \vDash \theta_{\alpha,\beta}(\overline{a})$, in which case, for the same α and every β, $\mathcal{A} \vDash \theta_{\alpha,\beta}(\overline{a})$. We conclude that $\mathcal{A} \vDash \phi(\overline{a})$.

Suppose that ψ transfers across elementary extensions. We will show that ψ is equivalent to Force_{ψ}. Let $\mathbb{A}$ be a fragment containing ψ. If $\mathcal{A} \vDash \psi(\overline{a})$, then because ψ transfers across elementary extensions, for any $\mathbb{A}$-generic $\mathcal{G} \geq \mathcal{A}$, $\mathcal{G} \vDash \psi(\overline{a})$. Consequently, $\mathcal{A} \Vdash^* \psi(\overline{a})$, so $\mathcal{A} \vDash \text{Force}_{\psi}(\overline{a})$. Suppose conversely that $\mathcal{A} \vDash \text{Force}_{\psi}(\overline{a})$. Then, $\mathcal{A} \Vdash^* \psi(\overline{a})$. Let $\mathcal{G} \geq \mathcal{A}$ be $\mathbb{A}$-generic. Then, $\mathcal{G} \vDash \psi(\overline{a})$, so because ψ transfers across elementary extensions, $\mathcal{A} \vDash \psi(\overline{a})$. $\square$

4.2.2 $\mathcal{L}_{\omega_1,\omega}$ and the Malitz Interpolation Theorem

Note that in Theorem 1.2, even if the formula $\psi(\bar{x})$ is in $\mathcal{L}_{\omega_1,\omega}$, the resulting formula in (2) may not be in $\mathcal{L}_{\omega_1,\omega}$ (as in Corollary 3.20). For the language $\mathcal{L}_{\omega_1,\omega}$, one can obtain a better result. First, we show what one can get from the Malitz interpolation theorem.

Theorem 4.5 (Malitz interpolation theorem [Mal69])**.** *Suppose the signature τ has no function symbols. Let φ, ψ be sentence of $\mathcal{L}_{\omega_1,\omega}$ such that ψ is universal ($\forall_1$), and $\varphi \vDash \psi$. Then, there is a universal sentence θ of $\mathcal{L}_{\omega_1,\omega}$ such that $\varphi \vDash \theta$, $\theta \vDash \psi$ and every symbol occurring in θ occurs in both φ and ψ.*

As a consequence of this, Malitz proves the following, which applies to any signature τ. (The version below, in which everything happens relative to a background theory represented by σ, appears in [Kei71]; Malitz also shows that if a formula is preserved both upwards and downwards then it is equivalent to a quantifier-free sentence, but this is not true relative to a sentence σ.)

Theorem 4.6 (Malitz [Mal69])**.** *Let φ and σ be sentences of $\mathcal{L}_{\omega_1,\omega}$. The following are equivalent.*

(1) If $\mathcal{A} \subset \mathcal{B}$, $\mathcal{A} \vDash \sigma$, $\mathcal{B} \vDash \sigma$, and $\mathcal{B} \vDash \varphi$, then $\mathcal{A} \vDash \varphi$.

(2) There is a universal sentence θ of $\mathcal{L}_{\omega_1,\omega}$ such that $\sigma \vDash \varphi \leftrightarrow \theta$.

Note that these theorems are valid only for $\mathcal{L}_{\omega_1,\omega}$ (and $\mathcal{L}_{\omega,\omega}$). Malitz [Mal71] has shown that the Craig interpolation theorem fails in $\mathcal{L}_{\kappa,\omega}$ for $\kappa > \omega_1$, and indeed there are examples with no interpolant in $\mathcal{L}_{\infty,\omega}$. We are not sure to what degree Theorem 4.6 fails in $\mathcal{L}_{\kappa,\omega}$, but Malitz [Mal69] has shown that there is a set of $\mathcal{L}_{\omega_1,\omega}$ sentences closed under substructures which is not equivalent to any set of universal $\mathcal{L}_{\omega_1,\omega}$ sentences. This set of sentences is, however, equivalent to a universal $\mathcal{L}_{\omega_2,\omega}$ sentence.

Using this theorem, we can give a different characterization of formulas $\mathcal{L}_{\omega_1,\omega}$ that are preserved by elementary extensions. We say that a formula is $\exists_1$ over finitary formulas if it can be obtained from finitary formulas by taking conjunctions, disjunctions, and existential quantification. Similarly, we say that a formula is $\forall_1$ over finitary formulas if it can be obtained from finitary formulas by taking conjunctions, disjunctions, and universal quantification.

Theorem 4.7. *Suppose ψ is a formula of $\mathcal{L}_{\omega_1,\omega}$. The following are equivalent.*

(1) Given $\mathcal{A} \leq \mathcal{B}$, $\mathcal{A} \models \varphi(\bar{a})$ if and only if $\mathcal{B} \models \varphi(\bar{a})$.

(2) There are formulas α and β of $\mathcal{L}_{\omega_1,\omega}$ such that α is $\forall_1$ over finitary formulas, β is $\exists_1$ over finitary formulas, and ψ, α, and β are all equivalent.

Proof Sketch. We omit the full proof as it is straightforward and this theorem will be subsumed by Theorem 1.3 to follow. Essentially one expands the signature by introducing a new relation symbol for each finitary formula and applies Theorem 4.6. □

Theorem 4.7 has the advantage that α and β are formulas of $\mathcal{L}_{\omega_1,\omega}$, while Theorem 1.2 has the advantage of giving us a single formula, all of whose quantifiers occur in finitary subformulas with complexity bounded by that of ψ. We can for the most part combine these advantages, in the following theorem. We say that a formula is quantifier-free over finitary formulas if it can be obtained from finitary formulas by taking conjunctions, disjunctions, and negation.

Theorem 1.3. *Let $\psi(\bar{x})$ be an infinitary $\mathcal{L}_{\omega_1,\omega}$ formula. The follow are equivalent:*

(1) Given $\mathcal{A} \leq \mathcal{B}$, $\mathcal{A} \models \psi(\bar{a})$ if and only if $\mathcal{B} \models \psi(\bar{a})$.

(2) $\psi(\overline{x})$ is equivalent to an $\mathcal{L}_{\omega_1,\omega}$ formula which is quantifier-free over finitary formulas.

Moreover, if these conditions hold and ψ is a $\forall_n$ formula (or an $\exists_n$ formula), then $\psi(\overline{x})$ is equivalent to an $\mathcal{L}_{\omega_1,\omega}$ formula which is quantifier-free over finitary $\exists_n/\forall_n$ formulas.

Proof. It is a straightforward induction that if ϕ is quantifier-free over finitary formulas (2), then ϕ transfers across elementary extensions (1).

Suppose ψ transfers across elementary extensions (1). Restricting the signature to symbols occurring in ψ, and adding constants for the free variables of ψ, it suffices to consider the case where the signature τ is countable and ψ is a sentence. Suppose that ψ is $\exists_n$. We say that $\mathcal{A} \equiv_n \mathcal{B}$ if $\mathcal{A}$ and $\mathcal{B}$ satisfy the same *finitary* $\exists_n$ sentences (and hence the same $\forall_n$ sentences).

First we show that whether $\mathcal{A} \models \psi$ depends only on the finitary n-theory of $\mathcal{A}$; that is, if $\mathcal{A} \equiv_n \mathcal{B}$, then $\mathcal{A} \models \psi$ if and only if $\mathcal{B} \models \psi$. By Theorem 1.2, ψ is equivalent to

$$\text{Force}_{\psi} = \bigvee_{\alpha} \bigwedge_{\beta} \theta_{\alpha,\beta}$$

where each $\theta_{\alpha,\beta}$ is a finitary $\exists_n$ sentence. If $\mathcal{A} \equiv_n \mathcal{B}$, then $\mathcal{A} \models \text{Force}_{\psi}$ if and only if $\mathcal{B} \models \text{Force}_{\psi}$, and so $\mathcal{A} \models \psi$ if and only if $\mathcal{B} \models \psi$.

Let D be the set of (finitary) $\exists_n$ sentences of $\mathcal{L}_{\omega,\omega}$. Given a set $T \subseteq D$, we identify T with the infinitary formula

$$\xi_T = \bigwedge_{\varphi \in T} \varphi \wedge \bigwedge_{\varphi \notin T} \neg\varphi.$$

Consider the set

$$X_{\psi} = \{T \subseteq D \mid \xi_T \text{ is satisfiable and } \xi_T \models \psi\} \subset 2^D$$

D is countably infinite, so we can identify 2^D with Cantor space, with subbasic clopen sets $[\theta] = \{T \subseteq D \mid \theta \in T\}$ and $[\neg\theta] = \{T \subseteq D \mid \theta \notin T\}$, for $\theta \in D$. We will show that X_ψ is a Borel set by showing that it is $\mathbf{\Sigma}^1_1$ and $\mathbf{\Pi}^1_1$. If there is a countable model $\mathcal{A} \models \xi_T \wedge \psi$, then ξ_T is satisfiable. For any $\mathcal{B} \models \xi_T$, $\mathcal{B} \equiv_n \mathcal{A}$, so $\mathcal{B} \models \psi$, by the above considerations. Thus, $\xi_T \models \psi$, so $T \in X_\psi$. Conversely, if $T \in X_\psi$, the Löwenheim-Skolem theorem for $\mathcal{L}_{\omega_1,\omega}$ implies that $\xi_T \wedge \psi$ has a countable model. Thus

$$X_\psi = \{T \subseteq D \mid \xi_T \wedge \psi \text{ has a countable model}\}$$

which is $\mathbf{\Sigma}^1_1$. On the other hand, the Löwenheim-Skolem theorem implies that if $\xi_T \not\models \psi$, there is a countable model $\mathcal{A} \models \xi_T \wedge \neg\psi$. Consequently,

$$X_\psi = \{T \subseteq D \mid \xi_T \wedge \neg\psi \text{ does not have a countable model }\} \cap \{T \subseteq D \mid \xi_T \text{ is satisfiable}\}$$

which is the intersection of a $\mathbf{\Pi}^1_1$ set with a Borel set, so is $\mathbf{\Pi}^1_1$.

We can assign to each Borel set $Y \subseteq 2^D$ an $\mathcal{L}_{\omega_1,\omega}$ sentence ϕ_Y which is quantifier-free over finitary $\exists_n/\forall_n$ formulas, and equivalent to the $\mathcal{L}_{(2^\omega)^+,\omega}$ sentence $\bigvee\!\!\!\bigvee_{T \in Y} \xi_T$.

- If $Y = [\theta]$ we can take $\phi_Y = \theta$, and if $Y = [\neg\theta]$ we can take $\phi_Y = \neg\theta$.
- If $Y = Z^C$, we can take $\phi_Y = \neg\phi_Z$. (This works because ξ_T and $\xi_{T'}$ are always inconsistent for $T \neq T'$.)
- If $Y = \bigcup_n Z_n$, we can take $\phi_Y = \bigvee\!\!\!\bigvee_n \phi_{Z_n}$.

Any Borel set can be built starting from subbasic clopen sets using countable unions and complements, so we can construct a ϕ_Y for any Y.

We will now show that ψ is equivalent to ϕ_{X_ψ}. If $\mathcal{A} \models \phi_{X_\psi}$, then for some $T \in X_\psi$, $\mathcal{A} \models \xi_T$. Because $\xi_T \models \psi$, $\mathcal{A} \models \psi$. Conversely, suppose $\mathcal{A} \models \psi$. Let $T = \{\theta \in D | \mathcal{A} \models \theta\}$ so that $\mathcal{A} \models \xi_T$. If $\mathcal{B} \models \xi_T$, then $\mathcal{A} \equiv_n \mathcal{B}$, so $\mathcal{B} \models \psi$. Therefore $\xi_T \models \psi$ and so $T \in X_\psi$. Because $\mathcal{A} \models \xi_T$, $\mathcal{A} \models \phi_{X_\psi}$.

Thus ψ is equivalent to the formula ϕ_{X_ψ} which is quantifier-free over finitary $\exists_n/\forall_n$ formulas. If we do not assume ψ to be an infinitary $\exists_n$ or $\forall_n$ formula, we have that if $\mathcal{A} \equiv \mathcal{B}$, $\mathcal{A} \models \psi$ if and only if $\mathcal{B} \models \psi$. Replacing D with all of $\mathcal{L}_{\omega,\omega}$, the rest of the argument goes through as before. In this case, ϕ_{X_ψ} is just quantifier-free over finitary formulas. □

References

[Kei71] H. Jerome Keisler. *Model theory for infinitary logic. Logic with countable conjunctions and finite quantifiers*. Studies in Logic and the Foundations of Mathematics, Vol. 62. North-Holland Publishing Co., Amsterdam-London, 1971.

[LE65] E. G. K. Lopez-Escobar. An interpolation theorem for denumerably long formulas. *Fund. Math.*, 57:253–272, 1965.

[Lou80] Alain Louveau. A separation theorem for Σ^1_1 sets. *Trans. Amer. Math. Soc.*, 260(2):363–378, 1980.

[Mal69] Jerome Malitz. Universal classes in infinitary languages. *Duke Math. J.*, 36:621–630, 1969.

[Mal71] Jerome Malitz. Infinitary analogs of theorems from first order model theory. *J. Symbolic Logic*, 36:216–228, 1971.

[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.