Notre Dame Journal of Formal Logic
Volume 49, Number 3, 2008

Quantifier Elimination for a Class of Intuitionistic Theories

Ben Ellison, Jonathan Fleischmann,
Dan McGinn, and Wim Ruitenburg

Abstract From classical, Fraïssé-homogeneous, ($\leq$ ω)-categorical theories over finite relational languages, we construct intuitionistic theories that are complete, prove negations of classical tautologies, and admit quantifier elimination. We also determine the intuitionistic universal fragments of these theories.

1 Introduction

It is often assumed that intuitionistic theories that admit quantifier elimination are nearly classical. We show that this is not the case. We present a straightforward method that converts a broad class of classical theories that admit quantifier elimination into intuitionistic ones.

Intuitionistic quantifier elimination has been studied before; see [11], [10], and [1], for example. Smoryński in [11] and Bagheri in [1] focus on intuitionistic theories that are in some ways nearly classical. Instead, we expand on the work in [10] and, in general, eliminate quantifiers in *very intuitionistic* theories, which in our case are theories that prove the negation of certain classical tautologies. Specifically, we start with a well-known class of classical theories over finite relational languages that admit quantifier elimination, are Fraïssé-homogeneous, and are ($\leq$ ω)-categorical. We call these theories JRS theories, after Jaśkowski, Rabin, and Scott, as explained in Section 2. We construct intuitionistic variations of the JRS theories and show these new theories retain the properties of completeness (Theorem 3.1) and quantifier elimination (Theorem 4.8), but in general are very intuitionistic. We show that if the morphism structure of the canonical Kripke model is sufficiently rich, then all formulas are equivalent to particularly simple quantifier-free formulas (Theorem 4.9). Our techniques for proving intuitionistic quantifier elimination are classical.

Received August 23, 2007; accepted March 7, 2008; printed June 6, 2008
2000 Mathematics Subject Classification: Primary, 03C10, 03F55; Secondary, 03C35, 03C90, 03F05
Keywords: intuitionistic predicate logic, quantifier elimination, Kripke model, Fraïssé homogeneous, normal forms
 10.1215/00294527-2008-012

In Section 5, as part of a deeper investigation into the idea of an intuitionistic model complete theory, we use the techniques and definitions of [6] to find the intuitionistic universal fragment of an intuitionistic JRS theory (Theorem 5.6). In the general intuitionistic case, quantifier-free formulas need not be universal formulas, in a sense that will be explained in Section 5. In our case, however, we show that all formulas are equivalent to quantifier-free, universal formulas (Theorem 5.2).

2 Classical JRS Theories

We review a special family of classical theories that admit quantifier elimination. We use the single turnstile $\vdash$ for "intuitionistically proves"; when we wish to indicate a classical proof, we use the $\vdash_c$ notation. Similarly, we write $\mathrm{Th}(\cdot)$ for the intuitionistic theory generated by a set of formulas or a structure, and $\mathrm{Th}_c(\cdot)$ for the classical theory. A theory Γ is *consistent* if $\bot \notin \Gamma$.

2.1 What is a JRS theory? We consider relational languages $\mathcal{L}$ that have only finitely many predicates $\{R_i\}_{i<r}$, all of positive arity. We use $\top$, $\bot$, $\wedge$, $\vee$, $\to$, $=$, $\exists$, and $\forall$ to form formulas of $\mathcal{L}$. Symbols $\top$ and $\bot$ are nullary logical operators as well as atoms. Negation $\neg\varphi$ is short for $\varphi \to \bot$.

Given a tuple $\mathbf{x} = x_0, x_1, \ldots, x_{n-1}$ of variables, the set $\mathcal{A}\mathrm{t}(\mathbf{x})$ of atoms with all free variables from $\mathbf{x}$ is finite. So the set $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$ of atoms and negated atoms over $\mathbf{x}$ is also finite. An $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-*type* is a subset $t \subseteq \mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$ such that its conjunction $\bigwedge t$, also written π_t or $\pi_t(\mathbf{x})$, is consistent. We write t^+ for the subcollection of atoms in t. We define formula π_t^+ to be the conjunction of atoms of t^+, and σ_t^- to be the disjunction of atoms whose negations occur in t. So $\pi_t \leftrightarrow (\pi_t^+ \wedge \neg\sigma_t^-)$ is a tautology. Formula π_t is called an $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-*description*. A maximal $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-type is called a *complete* $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-*type*, and its corresponding formula π_t a *complete* $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-*description*. Each atom of $\mathcal{A}\mathrm{t}(\mathbf{x})$ or its negation occurs in a complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-type. Given a model $\mathfrak{A}$ and $\mathbf{a} \in A$, $\mathbf{a}$ satisfies the complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-type $\mathrm{tp}_{\mathbf{a}} = (\mathrm{Th}_c(\mathfrak{A}) \cap \mathcal{A}\mathrm{t}^{\pm}(\mathbf{a}))[\mathbf{a}/\mathbf{x}]$, where $\mathrm{Th}_c(\mathfrak{A})$ is the theory of $\mathfrak{A}$ over the language $\mathcal{L}(A)$.

Suppose $n \geq 0$. Up to isomorphism, a complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-type t has a unique smallest model. Specifically, $\mathfrak{A}_t$ is the model formed from the variables $\{x_i\}_{i<n}$ by taking equivalence classes modulo the equivalence relation $x_i \sim x_j$ defined by $(x_i = x_j) \in t$. We write $\overline{x_i}$ or a_i for the equivalence class of x_i. Given $\mathbf{a} = a_0, \ldots, a_{n-1}$ and atom $\delta(\mathbf{x})$, set $\mathfrak{A}_t \models \delta(\mathbf{a})$ if and only if $\delta(\mathbf{x}) \in t$. So $\mathfrak{A}_t \models \pi_t(\mathbf{a})$. The size $|A_t|$ of model $\mathfrak{A}_t$ is called the *level* of t. We allow the empty structure.

Let u be an $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x}x_n)$-type. Define $d(u) = u \cap \mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$. Then $d(u)$ is an $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-type. If u is a complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x}x_n)$-type, then $d(u)$ is a complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-type. Given a complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x}x_n)$-type u, define δ_u to be the sentence

$$\forall \mathbf{x}(\pi_{d(u)} \to \exists x_n \pi_u).$$

We call such a sentence a JRS *sentence*. A (consistent) theory Γ over $\mathcal{L}$ is called a JRS *theory* if for all $\mathbf{x}x_n$ and complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x}x_n)$-types u that are consistent with Γ (that is, $\Gamma \cup \{\exists \mathbf{x}x_n \pi_u\}$ is consistent, or $\Gamma_\forall \nvdash \forall \mathbf{x}x_n \neg\pi_u$), we have $\delta_u \in \Gamma$.

As indicated by Bankston ([2], p. 962), this is not the first time that JRS theories and sentences have been studied. Gaifman attributes these sentences to Rabin and Scott (see [7], p. 15) while Lynch attributes them to Jaśkowski (see [9], p. 94), hence our choice of name.

2.2 Classical quantifier elimination The following are some well-known facts about JRS theories.

Theorem 2.1 *Let Γ be a* JRS *theory. Then, up to isomorphism, Γ has exactly one model of size $\leq \omega$. Additionally, this model is Fraïssé homogeneous; that is, isomorphisms between finite submodels extend to automorphisms.*

Proof The proof uses the axioms δ_u to complete a standard back and forth construction to extend finite isomorphisms to automorphisms. □

Recall that an existential formula is a *primitive formula* if its quantifier-free part is a conjunction of atoms and negated atoms.

Theorem 2.2 *Let Γ be a* JRS *theory, and let $\exists x_n \varphi(\mathbf{x}x_n)$ be a primitive formula. Then $\Gamma \vdash_c \exists x_n \varphi \leftrightarrow \bigvee_{s \in S} \pi_{d(s)}$, where $S = \{s : s$ is a complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x}x_n)$-type consistent with Γ and $\Gamma \vdash_c \pi_s \to \varphi\}$. In particular,* JRS *theories admit quantifier elimination.*

Proof Formula $\exists x_n \varphi$ is equivalent to $\bigvee_{s \in S} \exists x_n \pi_s$, where an empty disjunction is identified with $\bot$. Apply the JRS sentences of Γ: $\exists x_n \varphi$ is equivalent to $\bigvee_{s \in S} \pi_{d(s)}$. □

By the techniques in [8], Henson shows that there are continuum many JRS theories, even if the language has only one binary predicate. The work [2] of Bankston and one of the authors offers other construction techniques for JRS theories. Countable JRS theories can be built via certain types of games and can also be viewed as theories whose tree of finite substructures satisfies certain properties (see [2], Theorem 5.7). That is, given a theory Γ, form the following rooted tree T_Γ of types: for each $\mathbf{x} = x_0, \ldots, x_{n-1}$, take all complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-types of level n that are consistent with Γ (each such type essentially contains $\bigwedge_{i<j<n} x_i \neq x_j$). When we order these types by set inclusion, we get a tree with the minimal type $\{\top, \neg\bot\}$ as its root, and with finitely many nodes at each level. Obviously, T_Γ is uniquely determined by the universal fragment $\Gamma_\forall$ of Γ.

Given a universal theory Π, we define the JRS *extension* Γ of Π as the theory axiomatizable by Π and all JRS sentences δ_u for which $\Pi \nvdash \forall \mathbf{x} \neg \pi_u$. For a given universal theory Π, the consistency of the JRS extension is nicely expressible as a model-theoretic property on the collection of finite substructures $\mathfrak{A}_t$ of Π. A class of models K has the *amalgamation* property if for all models $\mathfrak{A}$, $\mathfrak{B}$, and $\mathfrak{C}$ in K where $\mathfrak{A}$ embeds in $\mathfrak{B}$ and $\mathfrak{A}$ embeds in $\mathfrak{C}$, there is a model $\mathfrak{D}$ in K such that $\mathfrak{B}$ embeds in $\mathfrak{D}$, $\mathfrak{C}$ embeds in $\mathfrak{D}$, and this diagram commutes. If K includes the empty structure, then the amalgamation property immediately implies the joint embedding property. This particularly applies to Theorem 2.3.

Theorem 2.3 *The* JRS *extension Γ of a universal theory Π is consistent if and only if the collection of models of the form $\mathfrak{A}_t$, for $t \in T_\Pi$, has the amalgamation property. If Γ is consistent, then $\Gamma_\forall = \Pi$.*

Proof First, suppose Γ is consistent. Let $\mathfrak{A}$ be the unique (up to isomorphism) model of Γ of size $\leq \omega$. Consider finite models $\mathfrak{A}_t$, $\mathfrak{A}_u$, and $\mathfrak{A}_v$ of $\Gamma_\forall$ and suppose that $\mathfrak{A}_t$ embeds in both $\mathfrak{A}_u$ and $\mathfrak{A}_v$. Without loss of generality, we may assume that u and v are complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x}x_n)$-types and that t is a complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-type. For some $\mathbf{a} \in A$, $\mathfrak{A}$ satisfies $\pi_t(\mathbf{a})$, δ_u and δ_v, so we have $\mathfrak{A} \models \exists x \pi_u(\mathbf{a}x) \wedge \exists x \pi_v(\mathbf{a}x)$. Fix $\mathbf{a}$, b

and c such that $\mathfrak{A} \models \pi_u(\mathbf{a}b) \wedge \pi_v(\mathbf{a}c)$. Let $w = \mathrm{tp}_{\mathbf{a}bc}$. Then $\mathfrak{A}_w$ is the amalgam of $\mathfrak{A}_u$ and $\mathfrak{A}_v$ over $\mathfrak{A}_t$.

Conversely, suppose that the collection of models of Π of the form $\mathfrak{A}_t$ has the amalgamation property. We sketch a construction of a model $\mathfrak{A}$ of Γ as the limit of an ω-chain of models of the form $\mathfrak{A}_t$. Suppose we have a model $\mathfrak{A}_t$ of size n. For each complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x}x_n)$-type u consistent with Π and for all $\mathbf{a} \in A_t$ such that $\mathfrak{A}_t \models \pi_{d(u)}(\mathbf{a})$ there is an amalgam $\mathfrak{A}_{(u,\mathbf{a})}$ of $\mathfrak{A}_t$ and $\mathfrak{A}_u$ over $\mathfrak{A}_{d(u)}$. As next model in the ω-chain, take the amalgam of all $\mathfrak{A}_{(u,\mathbf{a})}$ over $\mathfrak{A}_t$. So Γ is consistent.

For the last claim, it suffices to show that every finite structure of Π embeds into $\mathfrak{A}$, the unique largest model of size $\leq \omega$. Proceed by induction on the number of free variables in complete types consistent with Π. If u is a complete $\mathcal{A}\mathrm{t}^{\pm}$-type consistent with Π, then so is $d(u)$. By the inductive hypothesis, $\mathfrak{A}_{d(u)}$ embeds into $\mathfrak{A}$. By the JRS axiom δ_u, $\mathfrak{A}_u$ also embeds into $\mathfrak{A}$. □

2.3 Classical examples We present some examples of JRS theories and construction methods of new JRS theories from old.

Example 2.4 Let $\mathcal{L}$ be any language with finitely many predicate symbols of positive arity, and set Π to the minimal "empty" theory. Since all finite structures are allowed, amalgamation is obvious. By Theorem 2.3, the JRS extension of Π is consistent. This is an example of Burris's "theory of everything" [3].

Example 2.5 Let $\mathcal{L}$ be the minimal language (equality is the only relation). Theory $\Gamma = \Gamma_{\mathrm{e}}$ is the theory of infinite sets, with $\Gamma_\forall$ the "empty" theory. The tree T_Γ has just one node $t \supseteq \{x_i = x_j \to \bot \;:\; i < j < n\}$ at each level n.

Example 2.6 Let $\mathcal{L}$ be the language based on a new predicate $x \neq y$ for inequality. The theory of infinite sets $\Gamma = \Gamma_{\mathrm{ne}}$ has universal fragment axiomatizable by $x \neq y \leftrightarrow (x = y \to \bot)$. This direct translation makes Γ_{ne} "as JRS as" Γ_{e}.

Given a theory Γ, we write Γ_{UH} for the theory axiomatizable by its universal Horn fragment. Recall that models of Γ_{UH} are, up to isomorphism, submodels of products of models of Γ. If Γ is a JRS theory, then it is companionable with *few existential formulas*; that is, for each $\mathbf{x}$, there are only finitely many inequivalent (over Γ) existential formulas with variables from $\mathbf{x}$. So Γ_{UH} has a model companion $(\Gamma_{\mathrm{UH}})^*$ by Burris and Werner's work [4].

Example 2.7 It is a simple exercise to show that the theory of the random graph Γ_{g} is a JRS theory such that $(\Gamma_{\mathrm{g}})_{\mathrm{UH}} = (\Gamma_{\mathrm{ne}})_{\mathrm{UH}}$ (where we identify the single binary predicate R with the binary predicate $\neq$). Since Γ_{g} is model complete, $\Gamma_{\mathrm{g}} = ((\Gamma_{\mathrm{ne}})_{\mathrm{UH}})^*$. Comparing this with $\Gamma_{\mathrm{e}} = ((\Gamma_{\mathrm{e}})_{\mathrm{UH}})^*$ shows that seemingly trivial changes to language may significantly affect the derived universal Horn theories and their companions.

Example 2.8 Let $\mathcal{L}$ be the language based on $x \leq y$. The theory Γ_{lo} of dense linear order without endpoints is a well-known JRS theory.

Example 2.9 Let $\mathcal{L}$ be the language based on $x \leq y$. Let Γ_{p} be the theory of the random poset. Then it is a standard exercise to show $\Gamma_{\mathrm{p}} = ((\Gamma_{\mathrm{lo}})_{\mathrm{UH}})^*$ (see [5], p. 132, for example). Additionally, $\Gamma_{\mathrm{p}} = ((\Delta)_{\mathrm{UH}})^*$ where Δ is the non-JRS but obviously model complete trivial theory of a two-node linear order.

Note that $(\Gamma_{\mathrm{UH}})^*$ need not be a JRS theory, even if Γ is the JRS theory of a finite model.

3 Intuitionistic Theories from JRS Theories

Given a (classical) JRS theory Γ_{JRS} and its unique (up to isomorphism) model $\mathfrak{A}_{\text{JRS}}$ of size $\leq \omega$, we construct the Kripke model $\mathfrak{A}_{\text{M}}$ as follows. We follow notational conventions in [6]; our Kripke models are functors from small categories to the category of $\mathcal{L}$-structures and morphisms. The underlying category of $\mathfrak{A}_{\text{M}}$ consists of a single node with associated node structure $\mathfrak{A}_{\text{JRS}}$. We include all morphisms from $\mathfrak{A}_{\text{JRS}}$ to $\mathfrak{A}_{\text{JRS}}$ as arrows. Let Γ_{M} be the intuitionistic theory of $\mathfrak{A}_{\text{M}}$.

We can choose $\mathfrak{A}_{\text{M}}$ to be countable and get the same theory Γ_{M}. Instead of including all morphisms, let $\mathfrak{A}'_{\text{M}}$ have single node structure $\mathfrak{A}_{\text{JRS}}$ and include only a collection of morphisms closed under composition such that every finite graph of an endomorphism of $\mathfrak{A}_{\text{JRS}}$ has a complete endomorphism extension in the collection. A straightforward proof by induction on sentence complexity shows that $\mathfrak{A}_{\text{M}}$ and $\mathfrak{A}'_{\text{M}}$ have the same intuitionistic theory Γ_{M}. So our Kripke model can be chosen countable—take a category of countably many morphisms and a single countable object.

Theorem 3.1 *Γ_{M} is complete.*

Proof Let φ be an $\mathcal{L}$-sentence. If $\mathfrak{A}_{\text{M}} \Vdash \varphi$, then we are done. Otherwise, $\mathfrak{A}_{\text{M}} \nVdash \varphi$. But we have only one node, so $\mathfrak{A}_{\text{M}} \Vdash \neg\varphi$. □

Theorem 3.1 in no way implies that Γ_{M} proves classical logic. For example, if there is an endomorphism of $\mathfrak{A}_{\text{JRS}}$ which is not an embedding, then for some R_i and some $\mathbf{a}$ we have $\mathfrak{A}_{\text{M}} \nVdash R_i(\mathbf{a}) \vee \neg R_i(\mathbf{a})$, so $\mathfrak{A}_{\text{M}} \Vdash \neg\forall\mathbf{x}(R_i(\mathbf{x}) \vee \neg R_i(\mathbf{x}))$. In [10], Ruitenburg introduces one concept of a *very intuitionistic* theory to distinguish theories that are somehow even more "not classical." The two theories in [10], involving equality and linear order, are both examples of very intuitionistic theories. In general, suppose that instead of just one nonembedding endomorphism, we have two endomorphisms f and g, tuples $\mathbf{a}$ and $\mathbf{b}$, and formulas φ and ψ such that $\mathfrak{A}_{\text{M}} \Vdash \varphi(f\mathbf{a})$ and $\mathfrak{A}_{\text{M}} \nVdash \psi(f\mathbf{b})$, as well as $\mathfrak{A}_{\text{M}} \nVdash \varphi(g\mathbf{a})$ and $\mathfrak{A}_{\text{M}} \Vdash \psi(g\mathbf{b})$, as holds for the two examples from [10]. Then $\Gamma_{\text{M}} \vdash \neg\forall\mathbf{xy}((\varphi(\mathbf{x}) \to \psi(\mathbf{y})) \vee (\psi(\mathbf{y}) \to \varphi(\mathbf{x})))$, and therefore Γ_{M} is a very intuitionistic theory.

However, if $\mathfrak{A}_{\text{JRS}}$ is such that every endomorphism is also an embedding, then theory Γ_{M} is not of new interest to us, because of the following theorem.

Theorem 3.2 *If all endomorphisms of $\mathfrak{A}_{\text{JRS}}$ are embeddings, then $\Gamma_{\text{M}} = \Gamma_{\text{JRS}}$, and so Γ_{M} is a classical theory.*

Proof Since Γ_{JRS} admits quantifier elimination, it is model complete. Thus, all embeddings of Γ_{JRS} models are elementary embeddings. Apply Theorem A.1 in the Appendix. □

The examples from [10], as well as the examples from Subsection 2.3, satisfy the following special condition: We say that a model $\mathfrak{A}$ is *morphism homogeneous* if whenever $\mathbf{a}, \mathbf{b} \in A$ are such that $\text{tp}^+_{\mathbf{a}} \subseteq \text{tp}^+_{\mathbf{b}}$ then there is an endomorphism f of $\mathfrak{A}$ such that $f(\mathbf{a}) = \mathbf{b}$. A classical JRS theory Γ_{JRS} is *morphism homogeneous* if its unique countable model $\mathfrak{A}_{\text{JRS}}$ is. We show in Theorem 4.9 that if $\mathfrak{A}_{\text{JRS}}$ is morphism homogeneous, then Γ_{M} admits a particularly elegant kind of quantifier elimination.

Example 3.3 Not all $\mathfrak{A}_{\text{JRS}}$ are morphism homogeneous. Let $\mathcal{L}$ be the language with unary predicate $P(x)$ and binary predicate $x < y$, and let Γ_{JRS} be the (classical) theory of the finite model $\mathfrak{A}_{\text{JRS}}$ with domain $A_{\text{JRS}} = \{a, b\}$ such that

$\mathfrak{A}_{\mathrm{JRS}} \models \neg P(a) \wedge P(b) \wedge (a < b)$ and no other nontrivial atomic sentences. We have that $\mathrm{tp}_a^+ \subseteq \mathrm{tp}_b^+$ (in fact, $\mathrm{tp}_b^+ = \mathrm{tp}_a^+ \cup \{P(x)\}$). However, there is no morphism of $\mathfrak{A}_{\mathrm{JRS}}$ taking a to b. That is, assume f is a morphism such that $f(a) = b$. Then we must have $\mathfrak{A}_{\mathrm{JRS}} \models f(a) < f(b)$. But this is not true if $f(a) = b$, as $\mathfrak{A}_{\mathrm{JRS}} \models \forall x\, \neg(b < x)$.

4 Intuitionistic Quantifier Elimination in Γ_{M}

Recall that a theory has few (quantifier-free) formulas if for all $\mathbf{x} = x_0, x_1, \ldots, x_{n-1}$ there are finitely many nonequivalent (quantifier-free) formulas with all free variables from among $\mathbf{x}$. All classical theories over the finite relational language $\mathcal{L}$ have few quantifier-free formulas. So by quantifier elimination, Γ_{JRS} has few formulas. We show that the intuitionistic theory Γ_{M} admits quantifier elimination and also has few formulas. Our methods are classical.

Given a finite list of variables $\mathbf{x} = x_0, x_1, \ldots, x_{n-1}$, we first consider the complexity over Γ_{M} of the collection of quantifier-free formulas with all free variables from $\mathbf{x}$. Let $\mathfrak{C}(\mathbf{x})$ be the following Kripke model. As nodes for the underlying category $\mathbf{C}(\mathbf{x})$ we take all complete $\mathcal{A}\mathrm{t}^{\pm}(\mathbf{x})$-types t that are (classically) consistent with Γ_{JRS}. We turn $\mathbf{C}(\mathbf{x})$ into a poset category as follows. Given a pair of nodes t and u, we set $t \leq u$ exactly when there are $\mathbf{a} \in A_{\mathrm{JRS}}$ and endomorphism f of $\mathfrak{A}_{\mathrm{JRS}}$ such that $t = \mathrm{tp}_{\mathbf{a}}$ and $u = \mathrm{tp}_{f(\mathbf{a})}$ (that is, $\mathfrak{A}_{\mathrm{JRS}} \models \pi_t(\mathbf{a}) \wedge \pi_u(f(\mathbf{a}))$). So $t \leq u$ implies $t^+ \subseteq u^+$. To each node t we associate finite classical model $\mathfrak{A}_t$. If $t \leq u$, then the morphism sends the equivalence class $\overline{x_i}(t)$ of x_i in $\mathfrak{A}_t$ to the equivalence class $\overline{x_i}(u)$ of x_i in $\mathfrak{A}_u$. We write $\overline{x_i}$ for the "global" element $t \mapsto \overline{x_i}(t)$ of $\mathfrak{C}(\mathbf{x})$. The collection of nodes $|\mathbf{C}(\mathbf{x})|$ is finite. Note that $\mathfrak{A}_{\mathrm{JRS}}$ is morphism homogeneous exactly when $t^+ \subseteq u^+$ implies $t \leq u$ for every t and u in $|\mathbf{C}(\mathbf{x})|$.

Lemma 4.1 *Let $\varphi(\mathbf{x})$ be quantifier-free, and $\mathbf{a} \in A_{\mathrm{JRS}}$. Then $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(\mathbf{a})$ if and only if $\mathrm{tp}_{\mathbf{a}} \Vdash \varphi(\overline{\mathbf{x}}(\mathrm{tp}_{\mathbf{a}}))$.*

Proof We complete the proof by induction on the complexity of φ for all elements $\mathbf{a}$ simultaneously. The case for atoms and the induction steps for $\wedge$ and $\vee$ are easy. Let φ equal $\psi \to \theta$.

Suppose $\mathfrak{A}_{\mathrm{M}} \Vdash \psi(\mathbf{a}) \to \theta(\mathbf{a})$. Let $\mathrm{tp}_{\mathbf{a}} \leq u$ such that $u \Vdash \psi(\overline{\mathbf{x}}(u))$. It suffices to show that $u \Vdash \theta(\overline{\mathbf{x}}(u))$. There is an endomorphism f such that $u = \mathrm{tp}_{f(\mathbf{a})}$. By the inductive hypothesis, $\mathfrak{A}_{\mathrm{M}} \Vdash \psi(f(\mathbf{a}))$. By supposition, $\mathfrak{A}_{\mathrm{M}} \Vdash \theta(f(\mathbf{a}))$. So again by the inductive hypothesis, $u \Vdash \theta(\overline{\mathbf{x}}(u))$.

Conversely, suppose $\mathrm{tp}_{\mathbf{a}} \Vdash \psi(\overline{\mathbf{x}}(\mathrm{tp}_{\mathbf{a}})) \to \theta(\overline{\mathbf{x}}(\mathrm{tp}_{\mathbf{a}}))$. Let f be an endomorphism such that $\mathfrak{A}_{\mathrm{M}} \Vdash \psi(f(\mathbf{a}))$. It suffices to show $\mathfrak{A}_{\mathrm{M}} \Vdash \theta(f(\mathbf{a}))$. By the inductive hypothesis, $\mathrm{tp}_{f(\mathbf{a})} \Vdash \psi(\overline{\mathbf{x}}(\mathrm{tp}_{f(\mathbf{a})}))$. By definition $\mathrm{tp}_{\mathbf{a}} \leq \mathrm{tp}_{f(\mathbf{a})}$ so, by supposition, $\mathrm{tp}_{f(\mathbf{a})} \Vdash \theta(\overline{\mathbf{x}}(\mathrm{tp}_{f(\mathbf{a})}))$. Again by the inductive hypothesis, $\mathfrak{A}_{\mathrm{M}} \Vdash \theta(f(\mathbf{a}))$. □

To each quantifier-free $\varphi(\mathbf{x})$ assign $[\![\varphi(\overline{\mathbf{x}})]\!] = \{t \in |\mathbf{C}(\mathbf{x})| \;:\; t \Vdash \varphi(\overline{\mathbf{x}}(t))\}$. We can rewrite Lemma 4.1 above as $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(\mathbf{a})$ exactly when $\mathrm{tp}_{\mathbf{a}} \in [\![\varphi(\overline{\mathbf{x}})]\!]$. The sets $[\![\varphi(\overline{\mathbf{x}})]\!]$ form a finite Heyting algebra of upward closed subsets of the poset $\mathbf{C}(\mathbf{x})$ given by

$$[\![\varphi \wedge \psi]\!] = [\![\varphi]\!] \cap [\![\psi]\!],$$

$$[\![\varphi \vee \psi]\!] = [\![\varphi]\!] \cup [\![\psi]\!], \text{ and}$$

$$[\![\varphi]\!] \cap [\![\psi]\!] \subseteq [\![\theta]\!] \text{ if and only if } [\![\varphi]\!] \subseteq [\![\psi \to \theta]\!],$$

where we write $[\![\varphi]\!]$ as short for $[\![\varphi(\overline{\mathbf{x}})]\!]$, and so on. Subsets of the form $[\![\varphi(\overline{\mathbf{x}})]\!]$ are *definable*. Upward closed subsets of $\mathbf{C}(\mathbf{x})$ form the *open* subsets of the usual poset topology. So definable subsets are open. Below we show that open subsets are definable.

Lemma 4.2 *For all quantifier-free formulas $\varphi(\mathbf{x})$ and $\psi(\mathbf{x})$ we have $\Gamma_{\mathrm{M}} \vdash \forall\mathbf{x}(\varphi \to \psi)$ exactly when $[\![\varphi(\overline{\mathbf{x}})]\!] \subseteq [\![\psi(\overline{\mathbf{x}})]\!]$. Modulo provable equivalence over Γ_{M}, there are for each $\mathbf{x}$ only finitely many quantifier-free formulas with all free variables from $\mathbf{x}$.*

Proof Suppose that $\mathfrak{A}_{\mathrm{M}} \Vdash \forall\mathbf{x}(\varphi(\mathbf{x}) \to \psi(\mathbf{x}))$. Let $t \in [\![\varphi(\overline{\mathbf{x}})]\!]$. It suffices to show $t \in [\![\psi(\overline{\mathbf{x}})]\!]$. There is $\mathbf{a} \in A_{\mathrm{JRS}}$ such that $t = \mathrm{tp}_{\mathbf{a}}$. By Lemma 4.1, $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(\mathbf{a})$. By supposition, $\mathfrak{A}_{\mathrm{M}} \Vdash \psi(\mathbf{a})$. Again by Lemma 4.1, $\mathrm{tp}_{\mathbf{a}} \in [\![\psi(\overline{\mathbf{x}})]\!]$.

Conversely, suppose $[\![\varphi(\overline{\mathbf{x}})]\!] \subseteq [\![\psi(\overline{\mathbf{x}})]\!]$. Let $\mathbf{a} \in A_{\mathrm{JRS}}$ be such that $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(\mathbf{a})$. It suffices to show $\mathfrak{A}_{\mathrm{M}} \Vdash \psi(\mathbf{a})$. By Lemma 4.1, $\mathrm{tp}_{\mathbf{a}} \in [\![\varphi(\overline{\mathbf{x}})]\!]$. By supposition, $\mathrm{tp}_{\mathbf{a}} \in [\![\psi(\overline{\mathbf{x}})]\!]$. By Lemma 4.1 we get $\mathfrak{A}_{\mathrm{M}} \Vdash \psi(\mathbf{a})$.

So $[\![\varphi]\!] = [\![\psi]\!]$ exactly when $\Gamma_{\mathrm{M}} \vdash \forall\mathbf{x}(\varphi \leftrightarrow \psi)$. The second claim now follows, as $|\mathbf{C}(\mathbf{x})|$ is finite. □

Given $t \in |\mathbf{C}(\mathbf{x})|$, define $\hat{t} = \{u \in |\mathbf{C}(\mathbf{x})| \ : \ t \leq u\}$ and $\check{t} = \{u \in |\mathbf{C}(\mathbf{x})| \ : \ u \not\leq t\}$. So $\hat{t}$ is the smallest open subset containing t, and $\check{t}$ is the largest open subset not containing t. Clearly, $\hat{t} \subseteq [\![\pi_t^+(\overline{\mathbf{x}})]\!]$.

Lemma 4.3 *Let $t \in |\mathbf{C}(\mathbf{x})|$. Then $\check{t} = [\![\pi_t^+(\overline{\mathbf{x}}) \to \sigma_t^-(\overline{\mathbf{x}})]\!]$.*

Proof Suppose $s \leq t$. Then there are $\mathbf{a} \in A_{\mathrm{JRS}}$ and endomorphism f such that $s = \mathrm{tp}_{\mathbf{a}}$ and $t = \mathrm{tp}_{f(\mathbf{a})}$. So $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_t^+(f(\mathbf{a}))$ and $\mathfrak{A}_{\mathrm{M}} \nVdash \sigma_t^-(f(\mathbf{a}))$. So $\mathfrak{A}_{\mathrm{M}} \nVdash \pi_t^+(\mathbf{a}) \to \sigma_t^-(\mathbf{a})$. By Lemma 4.1, $s = \mathrm{tp}_{\mathbf{a}} \notin [\![\pi_t^+(\overline{\mathbf{x}}) \to \sigma_t^-(\overline{\mathbf{x}})]\!]$.

Conversely, suppose $s \not\leq t$. There is $\mathbf{a} \in A_{\mathrm{JRS}}$ such that $s = \mathrm{tp}_{\mathbf{a}}$. It suffices to show that $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_t^+(\mathbf{a}) \to \sigma_t^-(\mathbf{a})$. Let $s \leq u$ and let f be an endomorphism such that $u = \mathrm{tp}_{f(\mathbf{a})}$ and $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_t^+(f(\mathbf{a}))$. Then by supposition, $u \neq t$ and therefore there is an atomic formula δ such that $(\neg\delta) \in t$ and $\mathfrak{A}_{\mathrm{M}} \Vdash \delta(f(\mathbf{a}))$. So $\mathfrak{A}_{\mathrm{M}} \Vdash \sigma_t^-(f(\mathbf{a}))$. □

Let $t \in |\mathbf{C}(\mathbf{x})|$. We write ρ_t^- or $\rho_t^-(\mathbf{x})$ for

$$\textstyle\bigwedge_u(\pi_u^+ \to \sigma_u^-),$$

where $\bigwedge$ ranges over all u such that $t^+ \subseteq u^+$ but $t \not\leq u$. An empty conjunction is identified with $\top$. We write ρ_t^+ or $\rho_t^+(\mathbf{x})$ for $\pi_t^+ \wedge \rho_t^-$.

Lemma 4.4 *Let $t \in |\mathbf{C}(\mathbf{x})|$. Then $\hat{t} = [\![\rho_t^+(\overline{\mathbf{x}})]\!]$. So all open subsets of $\mathbf{C}(\mathbf{x})$ are definable.*

Proof To show $\hat{t} \subseteq [\![\rho_t^+(\overline{\mathbf{x}})]\!]$, it suffices to show $t \in [\![\rho_t^+(\overline{\mathbf{x}})]\!]$. Obviously, $t \in [\![\pi_t^+(\overline{\mathbf{x}})]\!]$. Let u be such that $t^+ \subseteq u^+$ and $t \not\leq u$. Then, by Lemma 4.3, $t \in [\![\pi_u^+(\overline{\mathbf{x}}) \to \sigma_u^-(\overline{\mathbf{x}})]\!]$. And thus $t \in [\![\rho_t^+(\overline{\mathbf{x}})]\!]$.

Conversely, suppose $v \in [\![\rho_t^+(\overline{\mathbf{x}})]\!]$. There is $\mathbf{a} \in A_{\mathrm{JRS}}$ such that $v = \mathrm{tp}_{\mathbf{a}}$. Then $\mathfrak{A}_{\mathrm{M}} \Vdash \rho_t^+(\mathbf{a})$. So $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_t^+(\mathbf{a})$ and $t^+ \subseteq \mathrm{tp}_{\mathbf{a}}^+$. Let u be such that $t^+ \subseteq u^+$ and $t \not\leq u$. Then $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_u^+(\mathbf{a}) \to \sigma_u^-(\mathbf{a})$. By Lemma 4.3, $\mathrm{tp}_{\mathbf{a}} \neq u$. Thus $t \leq \mathrm{tp}_{\mathbf{a}} = v$.

The second claim follows from the fact that all open sets are finite unions of sets $\hat{t}$. □

An open subset U is called *prime* if whenever U is the union $U = V \cup W$ of two open subsets, then $U = V$ or $U = W$. A prime open subset has *depth n* if there is a sequence of prime open subsets $U_0 \subseteq U_1 \subseteq \cdots \subseteq U_n$ such that $U_i \neq U_{i+1}$ for all i and $U_n = U$, but there is no longer sequence with these properties. So the empty subset has depth 0. The following is now obvious.

Lemma 4.5 *In* $\mathbf{C}(\mathbf{x})$*, each open subset equals a finite union of prime open subsets. A nonempty open subset is prime if and only if it is of the form* $\hat{t}$*, for some* $t \in |\mathbf{C}(\mathbf{x})|$.

Proof All open subsets in the poset topology are finite unions of sets of the form $\hat{t}$, so it suffices to prove that sets $\hat{t}$ are prime. This is immediate since $\hat{t} \subseteq U$ is equivalent to $t \in U$. □

Corollary 4.6 *Over* Γ_{M}*, every quantifier-free formula* φ *is equivalent to formula* $\bigvee\{\rho_t^+ : t \in [\![\varphi]\!]\}$.

Proof Immediate from Lemmas 4.5 and 4.4. □

Lemma 4.7 *For all formulas* $\varphi(\mathbf{x}x_n)$*, and for all* $t \in \mathbf{C}(\mathbf{x}x_n)$*,* Γ_{M} *includes the sentence,*

$$\forall \mathbf{x}x_n(\varphi \wedge \rho_t^+ \to (\sigma_t^- \vee \forall x_n(\rho_t^+ \to \varphi))).$$

Proof Fix φ, $t \in \mathbf{C}(\mathbf{x}x_n)$ and $\mathbf{a}, b \in A_{\mathrm{JRS}}$ and suppose $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(\mathbf{a}b) \wedge \rho_t^+(\mathbf{a}b)$. If $\mathfrak{A}_{\mathrm{M}} \Vdash \sigma_t^-(\mathbf{a}b)$ then we are done, so suppose not. Then $t = \mathrm{tp}_{\mathbf{a}b}$. We need to show that for arbitrary $c \in A_{\mathrm{JRS}}$ and endomorphism f, if $\mathfrak{A}_{\mathrm{M}} \Vdash \rho_t^+(f(\mathbf{a})c)$ then $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(f(\mathbf{a})c)$. Fix such an element c and endomorphism f. Then $\mathrm{tp}_{f(\mathbf{a})c} \in \hat{t}$ by Lemma 4.4. So $\mathrm{tp}_{\mathbf{a}b} \leq \mathrm{tp}_{f(\mathbf{a})c}$ and there is a morphism g such that $\mathrm{tp}_{g(\mathbf{a}b)} = \mathrm{tp}_{f(\mathbf{a})c}$. By the first supposition, $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(g(\mathbf{a}b))$. By Fraïssé homogeneity, there is an automorphism h such that $h(g(\mathbf{a}b)) = f(\mathbf{a})c$, so $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi(f(\mathbf{a})c)$. □

We are now ready to prove our main result.

Theorem 4.8 *Theory* Γ_{M} *admits quantifier elimination.*

Proof We eliminate quantifiers from formulas of the form $\varphi \wedge \theta$ where θ is quantifier-free (we recover all formulas by letting θ be $\top$). By Corollary 4.6, θ is equivalent to a formula of the form $\bigvee_{t \in S}\{\rho_t^+\}$ for some set $S \subseteq |\mathbf{C}(\mathbf{x})|$. Thus, each $\varphi \wedge \theta$ is equivalent to $\bigvee_{t \in S}\{\varphi \wedge \rho_t^+\}$. So it suffices to eliminate quantifiers from formulas of the form $\varphi \wedge \rho_t^+$, where $t \in S$. Fix such a formula, and proceed by induction on the depth of $[\![\rho_t^+]\!]$ and the number of free variables of φ.

Given $\varphi \wedge \rho_t^+$, if we have no free variables in φ, then by Theorem 3.1, $\varphi \wedge \rho_t^+$ is equivalent to a quantifier-free formula (namely, ρ_t^+ or $\perp$). Otherwise, apply Lemma 4.7. There are two cases.

In the first case, we get $\varphi \wedge \rho_t^+ \wedge \sigma_t^-$. As above, we use Corollary 4.6 to rewrite $\varphi \wedge (\rho_t^+ \wedge \sigma_t^-)$ as $\bigvee_{u \in R}(\varphi \wedge \rho_u^+)$ for some set $R \subseteq |\mathbf{C}(\mathbf{x})|$. Since $\bigvee_{u \in R} \rho_u^+ \to (\rho_t^+ \wedge \sigma_t^-)$, each ρ_u^+ implies ρ_t^+. By Lemma 4.2, for each $u \in R$, $[\![\rho_u^+]\!] \subseteq [\![\rho_t^+]\!]$. Likewise, since each ρ_u^+ implies σ_t^-, $[\![\rho_u^+]\!] \subseteq [\![\sigma_t^-]\!]$. By Lemma 4.5, each $[\![\rho_u^+]\!]$ is prime, and therefore $[\![\rho_u^+]\!] \subseteq [\![\delta]\!]$ for some atom δ found in σ_t^-. So $[\![\rho_u^+]\!] \neq [\![\rho_t^+]\!]$. By our inductive hypothesis on depth, each $\varphi \wedge \rho_u^+$ is equivalent to a quantifier-free formula, and therefore $\varphi \wedge \rho_t^+$ is equivalent to a quantifier-free formula.

In the second case, we get $\varphi \wedge \rho_t^+ \wedge \forall x_n(\rho_t^+ \to \varphi)$, which is equivalent to $\forall x_n(\rho_t^+ \to \varphi) \wedge \rho_t^+$. By the inductive hypothesis on free variables, this is equivalent to a quantifier-free formula. □

As a corollary we get the following.

Theorem 4.9 *Let $\varphi(\mathbf{x})$ be a formula. Over Γ_{M}, φ is equivalent to a disjunction of formulas ρ_t^+ with $t \in |\mathbf{C}(\mathbf{x})|$. If Γ_{JRS} is morphism homogeneous, then φ is equivalent to a disjunction of conjunctions of atoms π_t^+, with $t \in |\mathbf{C}(\mathbf{x})|$.*

Proof The first claim is immediate from Corollary 4.6 and Theorem 4.8. If Γ_{JRS} is morphism homogeneous, then for each t, $\Gamma_{\mathrm{M}} \vdash \pi_t^+ \leftrightarrow \rho_t^+$. So every quantifier-free formula φ is equivalent to $\bigvee\{\pi_t^+ : t \in [\![\varphi]\!]\}$, and therefore to a disjunction of conjunctions of atoms. □

As an illustration of Theorem 4.9 in the presence of morphism homogeneity, see the quantifier elimination results about the two theories in [10].

5 The Universal Fragment of Γ_{M}

Every classical model complete theory is uniquely determined by its universal fragment. Given the universal fragment, one can recover the model companion as the largest inductive theory preserving this universal fragment. As a start to a generalization of this process to intuitionistic theories, we find the universal fragments of our intuitionistic theories that admit quantifier elimination. We first need to explain what we mean by an intuitionistic universal sentence. The definition is motivated by Theorem 5.1 below; see also [6].

Recall that a Kripke model is essentially a functor $\mathfrak{A}$ from a (small) category $\mathbf{A}$ to the category of classical $\mathcal{L}$-structures and morphisms. That is, to each i in $|\mathbf{A}|$ we assign a classical structure $\mathfrak{A}_i$, and to each arrow $f : i \to j$ in $\mathbf{A}$ we assign a morphism $\mathfrak{A}f : \mathfrak{A}_i \to \mathfrak{A}_j$. We define "Kripke submodel" by $\mathfrak{A} \subseteq \mathfrak{B}$ if and only if $\mathbf{A} \subseteq \mathbf{B}$ as categories, and all morphisms and node structures of $\mathfrak{A}$ are restrictions of the corresponding morphisms and node structures of $\mathfrak{B}$. A sentence is *universal* if it can be built from the atoms using the operations $\wedge$, $\vee$, $\to$, and $\forall$, with the restriction that no implications or universal quantifications occur in negative places.

Theorem 5.1 *An intuitionistic theory Δ is axiomatizable by universal sentences if and only if its class of Kripke models is closed under Kripke submodels.*

Proof Immediate from [6], Theorem 4.1. □

Note that in the absence of Excluded Middle, not every quantifier-free formula is equivalent to a universal formula. Therefore, the following is an addition to Theorem 4.9.

Theorem 5.2 *Let $\varphi(\mathbf{x})$ be a formula. Over Γ_{M}, φ is equivalent to a quantifier-free universal formula.*

Proof This easily follows from Theorem 4.9 since each ρ_t^+ is a universal formula. □

Next, we axiomatize the universal fragment of Γ_{M}.

Lemma 5.3 *Let $t \in |\mathbf{C}(\mathbf{x})|$. Then Γ_{M} includes universal sentence $\forall \mathbf{x}(\pi_t^+ \to (\sigma_t^- \vee \rho_t^-))$.*

Proof Fix $\mathbf{a} \in A_{\mathrm{JRS}}$ and suppose that $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_t^+(\mathbf{a})$. If $\mathfrak{A}_{\mathrm{M}} \Vdash \sigma_t^-(\mathbf{a})$, we are done, so suppose $\mathfrak{A}_{\mathrm{M}} \nVdash \sigma_t^-(\mathbf{a})$. Then $t = \mathrm{tp}_{\mathbf{a}}$. Suppose we have endomorphism f and $u \in \mathbf{C}(\mathbf{x})$ such that $t^+ \subseteq u^+$, $t \nsubseteq u$, and $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_u^+(f(\mathbf{a}))$. Since $t \nsubseteq u$, $u \neq \mathrm{tp}_{f(\mathbf{a})}$. So $\mathfrak{A}_{\mathrm{M}} \Vdash \sigma_u^-(f(\mathbf{a}))$. □

Lemma 5.4 *Let $t \notin |\mathbf{C}(\mathbf{x})|$. Then Γ_{M} includes universal sentence $\forall \mathbf{x}(\pi_t^+ \rightarrow \sigma_t^-)$.*

Proof Fix $\mathbf{a} \in A_{\mathrm{JRS}}$ and suppose that $\mathfrak{A}_{\mathrm{M}} \Vdash \pi_t^+(\mathbf{a})$. Since $\mathfrak{A}_{\mathrm{JRS}} \not\models \pi_t(\mathbf{a})$, we have $\mathfrak{A}_{\mathrm{JRS}} \models \sigma_t^-(\mathbf{a})$. So $\mathfrak{A}_{\mathrm{M}} \Vdash \sigma_t^-(\mathbf{a})$. □

Note that the sentences $\forall \mathbf{x}(\pi_t^+ \rightarrow \sigma_t^-)$ from Lemma 5.4 axiomatize the universal fragment of the classical theory Γ_{JRS}. Since these sentences are geometric, the following well-known result applies.

Lemma 5.5 *Let $\mathfrak{B}$ be a Kripke model and φ a geometric sentence. Then $\mathfrak{B} \Vdash \varphi$ if and only if for each node $k \in |\mathbf{B}|$, node structure $\mathfrak{B}_k \models \varphi$.*

The schemas from Lemmas 5.3 and 5.4 suffice.

Theorem 5.6 *The axiom schemas*

$$\forall \mathbf{x}(\pi_t^+ \rightarrow \sigma_t^-) \quad \text{for all } \mathbf{x} \text{ and } t \notin |\mathbf{C}(\mathbf{x})|, \quad \text{and}$$
$$\forall \mathbf{x}(\pi_t^+ \rightarrow (\sigma_t^- \vee \rho_t^-)) \quad \text{for all } \mathbf{x} \text{ and } t \in |\mathbf{C}(\mathbf{x})|$$

together axiomatize the universal fragment of Γ_{M}.

Proof Let Δ be the set of all universal sentences described above. Let $\mathfrak{B} \Vdash \Delta$ be a Kripke model. By [12], Theorem 2.6.8, and because $\mathcal{L}$ is countable, we may suppose that $\mathbf{B}$ is a tree (poset) of height ω, and for all $i \in |\mathbf{B}|$ the domain of the node structure $\mathfrak{B}_i$ is at most countable. Let $r \in |\mathbf{B}|$ be the root of $\mathbf{B}$. We construct a rooted Kripke model $\mathfrak{D}$ with root r such that $\mathfrak{B} \subseteq \mathfrak{D}$ and $\mathfrak{D} \Vdash \Gamma_{\mathrm{M}}$.

First we construct an intermediate rooted Kripke model $\mathfrak{C}$ with $\mathbf{C} = \mathbf{B}$, $\mathfrak{B}_i \subseteq \mathfrak{C}_i \cong \mathfrak{A}_{\mathrm{JRS}}$ for every $i \in |\mathbf{C}|$, and $\mathfrak{C}f \upharpoonright B_i = \mathfrak{B}f$ for every $f : i \rightarrow j$ in $\mathbf{C}$. The construction is by induction on the height of $\mathbf{C}$. Let $\mathfrak{C}_r = \mathfrak{A}_{\mathrm{JRS}}$. By Lemmas 5.4 and 5.5, every node structure $\mathfrak{B}_i$ is a model of $(\Gamma_{\mathrm{JRS}})_\forall$. So up to isomorphism, $\mathfrak{B}_i \subseteq \mathfrak{A}_{\mathrm{JRS}}$ for every $i \in |\mathbf{B}|$. So without loss of generality, we may suppose that $\mathfrak{B}_r \subseteq \mathfrak{C}_r$. Now suppose that $\mathfrak{C}_i$ is defined for some $i \in |\mathbf{C}|$, with $\mathfrak{B}_i \subseteq \mathfrak{C}_i \cong \mathfrak{A}_{\mathrm{JRS}}$. Let $j \in |\mathbf{C}|$ be any immediate successor of i, and let $f : i \rightarrow j$ be the unique arrow from i to j in $\mathbf{C}$. Without loss of generality, we may suppose that $\mathfrak{B}_j \subseteq \mathfrak{C}_i$. We claim that there exists a $\mathfrak{C}_j \cong \mathfrak{A}_{\mathrm{JRS}}$ such that $\mathfrak{B}_j \subseteq \mathfrak{C}_j$, and a morphism $\mathfrak{C}f : \mathfrak{C}_i \rightarrow \mathfrak{C}_j$ such that $\mathfrak{C}f \upharpoonright B_i = \mathfrak{B}f$. Let $\mathcal{L}^*$ be the language $\mathcal{L}$ extended by a new function symbol f^*, and let $\Theta = \mathrm{Th}_c(\mathfrak{C}_i) \cup \{f^*(b) = \mathfrak{B}f(b) : b \in B_i\} \cup (\mathrm{Th}_c(\mathfrak{C}_i) \cap \mathcal{A}t(C_i))[c/f^*(c), c \in C_i]$, where $\mathrm{Th}_c(\mathfrak{C}_i)$ is the theory of the classical model $\mathfrak{C}_i$ over the language $\mathcal{L}(C_i)$. Let Θ_0 be any finite subset of Θ. Then

$$\Theta_0 \subseteq \mathrm{Th}_c(\mathfrak{C}_i) \cup \{f^*(b) = \mathfrak{B}f(b) : b \in \mathbf{b}\} \cup (\mathrm{Th}_c(\mathfrak{C}_i) \cap \mathcal{A}t(C_i))[c/f^*(c), c \in C_i],$$

for some finite $\mathbf{b} \subseteq B_i$. Obviously, $t = \mathrm{tp}_{\mathbf{b}}$ is consistent with Γ_{JRS}. Let $u = \mathrm{tp}_{\mathfrak{B}f(\mathbf{b})}$. Then, since $\mathfrak{B}f$ is a morphism, we have $t^+ \subseteq u^+$. Assume $t \nsubseteq u$. Then $\mathfrak{B} \Vdash \forall \mathbf{x}(\pi_t^+ \rightarrow (\sigma_t^- \vee (\pi_u^+ \rightarrow \sigma_u^-)))$. Since $i \Vdash^{\mathfrak{B}} \pi_t^+(\mathbf{b})$, we have $i \Vdash^{\mathfrak{B}} \sigma_t^-(\mathbf{b}) \vee (\pi_u^+(\mathbf{b}) \rightarrow \sigma_u^-(\mathbf{b}))$. Since $\mathfrak{B}_i \models \pi_t(\mathbf{b})$, we have $i \nVdash^{\mathfrak{B}} \delta(\mathbf{b})$, for every $\neg\delta \in t$. So $i \nVdash^{\mathfrak{B}} \sigma_t^-(\mathbf{b})$. So we must have $i \Vdash^{\mathfrak{B}} \pi_u^+(\mathbf{b}) \to \sigma_u^-(\mathbf{b})$. Since $\mathfrak{B}f$ is a morphism, we have $j \Vdash^{\mathfrak{B}} \pi_u^+(\mathfrak{B}f(\mathbf{b}))$. By the definition of forcing, $j \Vdash^{\mathfrak{B}} \sigma_u^-(\mathfrak{B}f(\mathbf{b}))$. So $j \Vdash^{\mathfrak{B}} \delta(\mathfrak{B}f(\mathbf{b}))$ for some $\neg\delta \in u$. So $\mathfrak{B}_j \models \delta(\mathfrak{B}f(\mathbf{b}))$ for some $\neg\delta \in u$. Contradiction. So $t \leq u$. So there is an endomorphism $f^* : \mathfrak{C}_i \to \mathfrak{C}_i$ such that $f^* \restriction \mathbf{b} = \mathfrak{B}f \restriction \mathbf{b}$. Let $\mathfrak{C}_i^*$ be the expansion of $\mathfrak{C}_i$ to $\mathcal{L}^*$ where f^* is interpreted as this endomorphism. Then $\mathfrak{C}_i^* \models \Theta_0$. So by compactness, Θ is consistent. Let $\mathfrak{C}_j^*$ be a countable model of Θ, and let $\mathfrak{C}_j$ be the $\mathcal{L}$-reduct of $\mathfrak{C}_j^*$. Then $\mathfrak{C}_i \preceq \mathfrak{C}_j$, and $f^* : \mathfrak{C}_i \to \mathfrak{C}_j$ is a morphism such that $f^* \restriction B_i = \mathfrak{B}f$. (Note that f^* is a total function on $\mathfrak{C}_j$, but it is only a morphism on $\mathfrak{C}_i \subseteq \mathfrak{C}_j$.) Set $\mathfrak{C}f = f^*$. Since $\mathfrak{A}_{\mathrm{JRS}}$ is the unique model of Γ_{JRS} of size less than or equal to ω, we have $\mathfrak{C}_j \cong \mathfrak{A}_{\mathrm{JRS}}$. So the claim is proven. This completes the construction of $\mathfrak{C}$. Clearly, $\mathfrak{B} \subseteq \mathfrak{C}$.

Let $\mathfrak{D}$ be the extension of $\mathfrak{C}$ generated by adding for each $i \in |\mathbf{C}|$ all possible morphisms from $\mathfrak{C}_i$ to itself. Then for all $\varphi \in \mathcal{L}(A_{\mathrm{JRS}})$ we have $\mathfrak{D} \Vdash \varphi$ if and only if $\mathfrak{A}_{\mathrm{M}} \Vdash \varphi$, by a straightforward induction on the complexity of φ. So $\mathfrak{D} \Vdash \Gamma_{\mathrm{M}}$. Also $\mathfrak{B} \subseteq \mathfrak{D}$. So by Theorem 5.1, $\mathfrak{B}$ forces the universal fragment of Γ_{M}. So Δ axiomatizes the universal fragment of Γ_{M}. □

Appendix A Kripke Models of Classical Logic

It is well known that Kripke models satisfy classical logic exactly when all morphisms between node structures are elementary embeddings. See [11], p. 110, for one direction. For the reader's convenience, we include a full proof. Recall that classical predicate logic CQC is axiomatizable over intuitionistic logic by the schema $\forall\mathbf{x}(\varphi(\mathbf{x}) \vee \neg\varphi(\mathbf{x}))$.

Theorem A.1 *Let $\mathfrak{A}$ be a Kripke model. Then the following are equivalent:*

1. *For all arrows $f : k \to m$ of $\mathbf{A}$, morphism $\mathfrak{A}(f)$ is an elementary embedding. That is, for all $\mathcal{L}(A_k)$ sentences $\varphi(\mathbf{a})$,*
   $\mathfrak{A}_k \models \varphi(\mathbf{a})$ *if and only if* $\mathfrak{A}_m \models \varphi(\mathbf{a})^f$.
2. *For all nodes $k \in |\mathbf{A}|$, and every sentence φ in $\mathcal{L}$,*
   CQC $\vdash_c \varphi$ *implies* $k \Vdash \varphi$.
3. *For every node k and for every sentence $\varphi(\mathbf{a})$ in $\mathcal{L}(A_k)$ we have*
   $\mathfrak{A}_k \models \varphi(\mathbf{a})$ *if and only if* $k \Vdash \varphi(\mathbf{a})$.

Proof

(2) ⇒ (3) We proceed by induction on the complexity of sentences. (3) holds for all atomic sentences, while the induction steps for existential statements, conjunctions, and disjunctions all follow directly from the definitions.

Given a node k, suppose $\mathfrak{A}_k \models \psi \to \theta$, where (3) holds for ψ and θ. If $\mathfrak{A}_k \models \psi$, then $\mathfrak{A}_k \models \theta$. By the inductive hypothesis, $k \Vdash \theta$, and so $k \Vdash \psi \to \theta$. Otherwise, $\mathfrak{A}_k \not\models \psi$. Then by the inductive hypothesis, $k \nVdash \psi$. By (2), $k \Vdash \psi \vee \neg\psi$, so $k \Vdash \neg\psi$. So $k \Vdash \psi \to \theta$.

Now suppose that $k \Vdash \psi \to \theta$, where (3) holds for ψ and θ. If $\mathfrak{A}_k \models \neg\psi$, then $\mathfrak{A}_k \models \psi \to \theta$ trivially. Otherwise, $\mathfrak{A}_k \models \psi$. By the inductive hypothesis, $k \Vdash \psi$, so $k \Vdash \theta$. By the inductive hypothesis again, $\mathfrak{A}_k \models \theta$. So $\mathfrak{A}_k \models \psi \to \theta$.

Suppose $\mathfrak{A}_k \models \forall x \psi(x)$, where (3) holds for $\psi(a)$, for all $a \in A_k$. Then, $\mathfrak{A}_k \models \psi(a)$ for all $a \in A_k$. By the inductive hypothesis, $k \Vdash \psi(a)$ for all $a \in A_k$. Assume $k \nVdash \forall x \psi(x)$. Then there exists $f : k \to m$ where $m \nVdash \psi^f(b)$, for some $b \in A_m$. By (2), $m \Vdash \psi^f(b) \vee \neg\psi^f(b)$, so $m \Vdash \neg\psi^f(b)$. Therefore, $m \Vdash \exists x \neg\psi^f(x)$. Now $k \Vdash \exists x \neg\psi(x)$ or $k \Vdash \neg\exists x \neg\psi(x)$ (again by (2)). The latter cannot hold, since $m \Vdash \exists x \neg\psi^f(x)$, so $k \Vdash \exists x \neg\psi(x)$. So, $k \Vdash \neg\psi(a)$ for some $a \in A_k$, a contradiction. Thus, $k \Vdash \forall x\, \psi(x)$.

Finally, suppose $k \Vdash \forall x\, \psi(x)$. So $k \Vdash \psi(a)$ for all $a \in A_k$. By the inductive hypothesis, $\mathfrak{A}_k \models \psi(a)$ for all $a \in A_k$. So $\mathfrak{A}_k \models \forall x\, \psi(x)$.

(3) $\Rightarrow$ (2) If CQC $\vdash_c \varphi$, then $\mathfrak{B} \models \varphi$ for all classical models $\mathfrak{B}$. Thus, given a node k, and a sentence φ proven by CQC, we have $\mathfrak{A}_k \models \varphi$. By (3), $k \Vdash \varphi$, proving (2).

(3) $\Rightarrow$ (1) Let $f : k \to m$, and suppose $\mathfrak{A}_k \models \varphi(\mathbf{a})$. By (3), $k \Vdash \varphi(\mathbf{a})$, and so $m \Vdash \varphi(\mathbf{a})^f$. By (3) again, $\mathfrak{A}_m \models \varphi(\mathbf{a})^f$.

(1) $\Rightarrow$ (3) We again proceed by induction on the complexity of sentences. By the definition of forcing, (3) always holds for atomic sentences, and the inductive steps for conjunctions, disjunctions, and existential statements are easy.

Suppose $\mathfrak{A}_k \models \psi \to \theta$. Let $f : k \to m$ be a morphism such that $m \Vdash \psi^f$. By the inductive hypothesis, $\mathfrak{A}_m \models \psi^f$. By (1), $\mathfrak{A}_m \models \psi^f \to \theta^f$, hence $\mathfrak{A}_m \models \theta^f$. By the inductive hypothesis, $m \Vdash \theta^f$, so $k \Vdash \psi \to \theta$.

Suppose $k \Vdash \psi \to \theta$. If $\mathfrak{A}_k \models \psi$ then, by the inductive hypothesis, $k \Vdash \psi$. Then $k \Vdash \theta$, so by the inductive hypothesis again, $\mathfrak{A}_k \models \theta$. Thus, $\mathfrak{A}_k \models \psi \to \theta$.

Suppose $\mathfrak{A}_k \models \forall x\, \psi(x)$, with (3) holding for $\psi^f(b)$, for all $b \in A_m$, where m is a node with morphism $f : k \to m$. Given such an f, by (1) we have $\mathfrak{A}_m \models \forall x\, \psi^f(x)$. Then, for all $a \in A_m$, $\mathfrak{A}_m \models \psi^f(a)$. By the inductive hypothesis, for every $a \in A_m$ we have $m \Vdash \psi^f(a)$. As f is arbitrary, we have that $k \Vdash \forall x\, \psi(x)$.

Finally, suppose $k \Vdash \forall x\, \psi(x)$. Then for all $a \in A_k$ we have $k \Vdash \psi(a)$. By the inductive hypothesis, $\mathfrak{A}_k \models \psi(a)$ for all $a \in A_k$. So $\mathfrak{A}_k \models \forall x\, \psi(x)$. □

References

[1] Bagheri, S. M., "Categoricity and quantifier elimination for intuitionistic theories," pp. 23–41 in *Logic in Tehran*, vol. 26 of *Lecture Notes in Logic*, Association for Symbolic Logic, La Jolla, 2006. Zbl 1106.03034. MR 2262312. 281

[2] Bankston, P., and W. Ruitenburg, "Notions of relative ubiquity for invariant sets of relational structures," *The Journal of Symbolic Logic*, vol. 55 (1990), pp. 948–86. Zbl 0726.03025. MR 1071308. 282, 283

[3] Burris, S., "The model completion of the class of $\mathcal{L}$-structures," *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, vol. 33 (1987), pp. 313–14. Zbl 0608.03011. MR 913293. 284

[4] Burris, S., and H. Werner, "Sheaf constructions and their elementary properties," *Transactions of the American Mathematical Society*, vol. 248 (1979), pp. 269–309. Zbl 0411.03022. MR 522263. 284

[5] van Dalen, D., *Logic and Structure*, 4th edition, Universitext. Springer, New York, 2004. Zbl 1048.03001. MR 2052431. 284

[6] Ellison, B., J. Fleischmann, D. McGinn, and W. Ruitenburg, "Kripke submodels and universal sentences," *Mathematical Logic Quarterly*, vol. 53 (2007), pp. 311–20. Zbl 1123.03008. MR 2330601. 282, 285, 289

[7] Gaifman, H., "Concerning measures in first order calculi," *Israel Journal of Mathematics*, vol. 2 (1964), pp. 1–18. Zbl 0192.03302. MR 0175755. 282

[8] Henson, C. W., "Countable homogeneous relational structures and $\aleph_0$-categorical theories," *The Journal of Symbolic Logic*, vol. 37 (1972), pp. 494–500. Zbl 0259.02040. MR 0321727. 283

[9] Lynch, J. F., "Almost sure theories," *Annals of Mathematical Logic*, vol. 18 (1980), pp. 91–135. Zbl 0433.03020. MR 578540. 282

[10] Ruitenburg, W., "Very intuitionistic theories and quantifier elimination," *The Review of Modern Logic*, vol. 10 (2004/05), pp. 99–112. MR 2238775. 281, 285, 289

[11] Smoryński, C., "Elementary intuitionistic theories," *The Journal of Symbolic Logic*, vol. 38 (1973), pp. 102–34. Zbl 0261.02033. MR 0327500. 281, 291

[12] Troelstra, A. S., and D. van Dalen, *Constructivism in Mathematics. An Introduction. Vol. 1*, vol. 121 of *Studies in Logic and the Foundations of Mathematics*, North-Holland Publishing Co., Amsterdam, 1988. Zbl 0653.03040. MR 966421. 290

Acknowledgments

The authors thank Asher Kach for his helpful suggestions.

Department of Mathematics
University of Wisconsin-Madison
480 Lincoln Drive
Madison WI 53706
ellison@math.wisc.edu

Department of Mathematics, Statistics and Computer Science
Marquette University
PO Box 1881
Milwaukee WI 53201
jonathan.fleischmann@marquette.edu

*Department of Mathematics
University of Wisconsin-Madison
480 Lincoln Drive
Madison WI 53706
mcginn@math.wisc.edu (corresponding author)

Department of Mathematics, Statistics and Computer Science
Marquette University
PO Box 1881
Milwaukee WI 53201
wim.ruitenburg@marquette.edu

*The corresponding author for this paper is Dan McGinn