

ON THE SCOPE OF THE UNIVERSAL-ALGEBRAIC APPROACH TO CONSTRAINT SATISFACTION*

MANUEL BODIRSKY^a, MARTIN HILS^b, AND BARNABY MARTIN^c

^a CNRS/LIX, École Polytechnique, France
e-mail address: bodirsky@lix.polytechnique.fr

^b Institut de Mathématiques de Jussieu, Université Paris Diderot - Paris 7, France
e-mail address: hils@logique.jussieu.fr

^c Engineering and Computing Sciences, Durham University, U.K.
e-mail address: barnabymartin@gmail.com

ABSTRACT. The universal-algebraic approach has proved a powerful tool in the study of the computational complexity of constraint satisfaction problems (CSPs). This approach has previously been applied to the study of CSPs with finite or (infinite) ω -categorical templates, and relies on two facts. The first is that in finite or ω -categorical structures \mathfrak{A} , a relation is primitive positive definable if and only if it is preserved by the polymorphisms of \mathfrak{A} . The second is that every finite or ω -categorical structure is homomorphically equivalent to a core structure.

In this paper, we present generalizations of these facts to infinite structures that are not necessarily ω -categorical. Specifically, we prove that every CSP can be formulated with a template \mathfrak{A} such that a relation is primitive positive definable in \mathfrak{A} if and only if it is first-order definable in \mathfrak{A} and preserved by the *infinitary* polymorphisms of \mathfrak{A} . Using existential positive closure we rederive and extend known results about cores, presenting the new notions of core theories (models of which will be cores) and core companions (which are defined analogously to model companions in model theory). We prove a uniqueness result for core companions that yields the uniqueness of model-complete cores of ω -categorical structures. Existential positive closure is also the crucial concept to give an exact characterization of those CSPs that can be formulated with (a finite or) an ω -categorical template.

Finally, we present applications of our general results to the description and analysis of the computational complexity of CSPs. In particular, we give general hardness criteria based on the absence of polymorphisms that depend on more than one argument, and we present a polymorphism-based description of those CSPs that are first-order definable (and therefore can be solved in polynomial time).

1998 ACM Subject Classification: F.4.1.

Key words and phrases: Constraint satisfaction; Model theory; Universal-algebraic approach.

* Many of the results of this paper have appeared in the extended abstract [7].

The first author has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 257039). The third author was supported by EPSRC grant EP/G020604.

1. INTRODUCTION

For a relational structure \mathfrak{A} over a finite signature the *constraint satisfaction problem* $\text{CSP}(\mathfrak{A})$ is the computational problem to decide whether a primitive positive first-order sentence φ – that is, the existential quantification of a conjunction of atomic formulas – is true on \mathfrak{A} . The case where the template \mathfrak{A} is finite has been extensively studied in the literature, and is known to comprise a significant microcosm of the complexity class NP (see, e.g., [24]). The universal-algebraic approach, of studying the invariance properties of relations under the action of polymorphisms, has been particularly powerful in the complexity analysis of finite-domain CSPs (see [23] as a starting point). This approach has also been successfully used in infinite-domain CSPs where the template is ω -categorical, i.e., is the unique countably infinite model of its first-order (fo) theory up to isomorphism – see, e.g., [10].

Many interesting problems can be formulated as infinite CSPs whose template is not ω -categorical. To illustrate the wealth of the class of CSPs studied in this paper, we present three concrete computational problems that can be formulated as $\text{CSP}(\mathfrak{A})$, for an infinite \mathfrak{A} . Each of these problems is solvable in polynomial time – and the proofs of this are generally non-trivial. The templates $(\mathbb{Z}; +, 1)$ and $(\mathbb{R}; +, 1)$, where $+$ is read as the ternary relation $x + y = z$, correspond to the solving of Linear Diophantine Equations and Linear Real Equations, respectively. Another template of interest relates to the Unification Problem. Let $\sigma := (f_1, f_2, \dots)$ be a functional signature, we form the template $(T; F_1, F_2, \dots)$, where T is the term algebra on σ built over a countably infinite set of variables, and each F_i is the relational form $f_i(t_1, \dots, t_{r_i}) = t_0$ of f_i over T .

Primitive positive definability of formulas. A central role is played in CSPs by the notion of primitive positive (pp) definability. Two templates give rise to the same CSP iff they agree on all pp-sentences (have the same pp-theory). A less central, but often equivalent role is played by existential positive (ep) definability. For example, two templates have the same pp-theory iff they have the same ep-theory. If $\langle \mathfrak{A} \rangle_{\text{pp}}$ are the relations pp-definable on \mathfrak{A} , then $\langle \mathfrak{A} \rangle_{\text{pp}} \subseteq \langle \mathfrak{A}' \rangle_{\text{pp}}$ implies that there is a polynomial time reduction from $\text{CSP}(\mathfrak{A}')$ to $\text{CSP}(\mathfrak{A})$. In the case that \mathfrak{A} is finite or ω -categorical, the set of relations $\langle \mathfrak{A} \rangle_{\text{pp}}$ are precisely those that are invariant under the polymorphisms of \mathfrak{A} – $\text{Inv}(\text{Pol}(\mathfrak{A}))$. We note that this relationship holds on some infinite structures which are not ω -categorical (an example is given in [29], also the natural numbers with the binary successor relation may easily be verified to have this property). In general, a CSP can not be formulated with (a finite or) an ω -categorical template (we return to this below). However, we are able to prove the existence of an equivalent template – with the same pp-theory – which enjoys some of the benign properties of finiteness or ω -categoricity. Given any \mathfrak{A} , we prove the existence of a highly saturated elementary extension \mathfrak{M} such that a relation is pp-definable on \mathfrak{M} iff it is fo-definable on \mathfrak{M} and invariant under the polymorphisms of (countably) infinite arity of \mathfrak{M} . In fact, we prove that this relationship holds for all saturated structures of cardinality at least 2^ω . But our construction obviates the need for the set-theoretic assumptions usually required to assert the existence of a saturated elementary extension of an arbitrary structure. However, in many concrete cases, such as for structures that are uncountably categorical, such saturated models arise directly. We go on to prove that each of the three assumptions – high saturation, infinitary (and not finitary) polymorphism and fo-intersection – is necessary. That is, we exhibit structures for which any two of these is insufficient for the respective connection.

Existential-positively closed models. We use existential-positively closed (epc) models to easily rederive and extend known results in the literature from [3]. In the case of finite-domain structures there is the well-established notion of core – a structure, homomorphically equivalent to the original, all of whose endomorphisms are embeddings (equivalently, in the finite, automorphisms). The core is unique up to isomorphism, and the task of problem classification is simplified as one may restrict attention to cores. For ω -categorical structures we no longer have uniqueness of cores in general, but it was shown in [3] that every ω -categorical structure is homomorphically equivalent to a model-complete core, i.e. to a structure whose endomorphisms preserve all first-order formulas, that is unique up to isomorphism and again ω -categorical. We present a new proof here, on the way introducing the concepts of core theories, models of which will be cores, and core companions. If an ep-theory S is contained in a model-complete core theory T – with the same existential positive and universal negative restriction – which will be called the *core companion* of S , then T is unique up to equivalence of theories. In the special case of ω -categorical structures, this yields the uniqueness result for model-complete cores. We give a necessary and sufficient condition as to when $\text{CSP}(\mathfrak{A})$ can be formulated with a template which has a model-complete core theory, based on whether the class of epc models of a restricted part of \mathfrak{A} 's $\forall\exists$ -theory is axiomatizable.

When distinct templates give rise to the same CSP, it might be the case that one is better behaved than another. For example, $(\mathbb{Z}; <)$ and $(\mathbb{Q}; <)$ share the same ep-theory (pp-theory); yet while $\text{Inv}(\text{Pol}(\mathbb{Z}; <)) \neq \langle (\mathbb{Z}; <) \rangle_{\text{pp}}$, we have $\text{Inv}(\text{Pol}(\mathbb{Q}; <)) = \langle (\mathbb{Q}; <) \rangle_{\text{pp}}$, as $(\mathbb{Q}; <)$ is ω -categorical. In the present paper we give a necessary and sufficient condition that a template \mathfrak{A} has an equivalent \mathfrak{A}' that is finite or ω -categorical – and consequently satisfies $\text{Inv}(\text{Pol}(\mathfrak{A}')) = \langle \mathfrak{A}' \rangle_{\text{pp}}$. The condition is that the number of maximal ep- n -types consistent with the theory of \mathfrak{A} is finite, for all n . It follows that none of the three examples of the previous page may be formulated with an ω -categorical template. Our result uses epc models and the direct limit construction.

Applications. We go on to consider the repercussions of our restricted relationship for the complexity of CSPs. Firstly, we show that existential positive (ep) and pp-definability coincide on a structure \mathfrak{A} iff all ω -polymorphisms of all elementary extensions of \mathfrak{A} are essentially unary. We demonstrate that the move to elementary extensions is necessary by giving a structure whose ω -polymorphisms include only projections but for which $(x = y \vee u = v)$ is not pp-definable. Using the notion of local refutability in [6], we note that if a structure \mathfrak{A} is not locally refutable, and all ω -polymorphisms of all elementary extensions of \mathfrak{A} are essentially unary, then $\text{CSP}(\mathfrak{A})$ is NP-hard.

Secondly, and introducing our philosophy to the work of [26], we present a polymorphism based description of those CSPs that are fo-definable. We show that $\text{CSP}(\mathfrak{A})$ is fo-definable if and only if \mathfrak{A} has an elementary extension which has a 1-tolerant polymorphism. It follows that such CSPs are polynomial-time solvable.

Thirdly, we recall a known relationship between certain binary injective polymorphisms and Horn definability (given in the context of ω -categorical structures in [4]). Considering as a polymorphism an embedding e of $(\mathbb{R}; +, 1)^2$ into $(\mathbb{R}; +, 1)$, we show that the recent complexity classification of [8] may be given a natural algebraic specification. The presence of the polymorphism e separates those fo-expansions of $(\mathbb{R}; +, 1)$ whose CSP is in P from those whose CSP is NP-complete. Thus we demonstrate that the presence of certain polymorphisms can delineate complexity even outside of the realm of ω -categoricity.

Related work. There are several extant works on notions of pp-definability over infinite structures, including those involving infinitary polymorphisms and infinitary relations [34, 27, 30]. Relational operations transcending normal pp-definitions are usually permitted, such as: infinite conjunction, infinite projection and various forms of monotone disjunction. In order for our results to be applicable to the (finite!) instances of CSPs, we are not able to sacrifice anything on the relational side, and so pp-definability must remain in its most basic form. This represents the principle difference between our work and those that have come before. We note that this is the first time that infinitary polymorphisms have been considered in connection with the complexity of CSPs.

An extended abstract of this paper appeared at conference as [7]. Section 4.2, as well as many previously omitted proofs, appear here for the first time.

2. PRELIMINARIES

2.1. Models, operations and theories. A *relational signature* (with constants) τ is a set of relation symbols R_i , each of which has an associated finite arity k_i , and constants c_i . We consider only relational signatures (with constants) in this paper. A (relational) *structure* \mathfrak{A} over the signature τ (also called τ -*structure*) consists of a set A (the *domain*) together with a relation $R^{\mathfrak{A}} \subseteq A^k$ for each relation symbol R of arity k from τ and a constant $c^{\mathfrak{A}} \in A$ for each constant symbol c .

Let \mathfrak{A} be a τ -structure, and let \mathfrak{A}' be a τ' -structure with $\tau \subseteq \tau'$. If \mathfrak{A} and \mathfrak{A}' have the same domain and $R^{\mathfrak{A}} = R^{\mathfrak{A}'}$, $c^{\mathfrak{A}} = c^{\mathfrak{A}'}$ for all $R, c \in \tau$, then \mathfrak{A} is called the τ -*reduct* (or simply *reduct*) of \mathfrak{A}' , and \mathfrak{A}' is called a τ' -*expansion* (or simply *expansion*) of \mathfrak{A} . If \mathfrak{A} is a τ -structure and $\langle a_\alpha \rangle_{\alpha < \beta}$ is a sequence of elements of A , then $(\mathfrak{A}; \langle a_\alpha \rangle_{\alpha < \beta})$ is the natural $\tau \cup \{c_\alpha : \alpha < \beta\}$ -expansion of \mathfrak{A} with β new constants, where c_α is interpreted by a_α . \mathfrak{A} is an *extension* of \mathfrak{B} , denoted $\mathfrak{B} \subseteq \mathfrak{A}$, if $B \subseteq A$ and for each R in τ , and for all tuples \bar{b} from B , $\bar{b} \in R^{\mathfrak{B}}$ iff $\bar{b} \in R^{\mathfrak{A}}$, and for each c in τ , $c^{\mathfrak{B}} = c^{\mathfrak{A}}$. Let $\langle b_\alpha \rangle_{\alpha < |B|}$ well-order the elements of \mathfrak{B} . \mathfrak{A} is an *elementary extension* of \mathfrak{B} , denoted $\mathfrak{B} \preceq \mathfrak{A}$, if it is an extension and, for each first-order (fo) $\tau \cup \{c_\alpha : \alpha < |B|\}$ -sentence φ , $(\mathfrak{B}, \langle b_\alpha \rangle_{\alpha < |B|}) \models \varphi$ iff $(\mathfrak{A}, \langle b_\alpha \rangle_{\alpha < |B|}) \models \varphi$.

An fo-formula is *existential positive* (ep) if it involves only existential quantification, conjunction and disjunction (no negation). Furthermore, if it involves no instances of disjunction, then it is termed *primitive positive* (pp). Note that we consider the boolean false \perp to be a pp-formula, and we always allow equalities in pp-formulas. Dual to ep are *universal negative* formulas, involving only universal quantification, conjunction and disjunction, with all atoms negated. Suppose \mathfrak{A} is a finite structure over a finite signature with domain $A := \{a_1, \dots, a_s\}$. Let $\theta(x_1, \dots, x_s)$ be the conjunction of the positive facts of \mathfrak{A} , where the variables x_1, \dots, x_s correspond to the elements a_1, \dots, a_s . That is, $R(x_{\lambda_1}, \dots, x_{\lambda_k})$ appears as an atom in θ iff $(a_{\lambda_1}, \dots, a_{\lambda_k}) \in R^{\mathfrak{A}}$. Define the *canonical query* $\varphi[\mathfrak{A}]$ of \mathfrak{A} to be the pp-formula $\exists x_1 \dots x_s. \theta(x_1, \dots, x_s)$. A set of formulas $\Phi := \Phi(x_1, \dots, x_n)$ with free variables x_1, \dots, x_n is called *satisfiable in* \mathfrak{A} if there are elements a_1, \dots, a_n from A such that for all sentences $\varphi \in \Phi$ we have $\mathfrak{A} \models \varphi(a_1, \dots, a_n)$. We say that Φ is *satisfiable* if there exists a structure \mathfrak{A} such that Φ is satisfiable in \mathfrak{A} .

A τ -*theory* is a set of τ -sentences; two theories are *equivalent* if they share the same models. For a τ -structure \mathfrak{A} , define the *theory* of \mathfrak{A} , $\text{Th}(\mathfrak{A})$, to be the set of τ -sentences true on \mathfrak{A} . Note that $\mathfrak{A} \preceq \mathfrak{B}$ implies that $\text{Th}(\mathfrak{A}) = \text{Th}(\mathfrak{B})$. Define $\text{Th}_{\forall^-}(\mathfrak{A})$ to be the set of universal negative sentences true on \mathfrak{A} .

For $n \geq 0$, an n -type of a theory T is a set $p := p(x_1, \dots, x_n)$ of formulas in the free variables x_1, \dots, x_n such that $p \cup T$ is satisfiable. In a similar manner, a *primitive positive* n -type (pp- n -type) of a theory T is a set of pp-formulas such that $p \cup T$ is satisfiable. A pp- n -type p of T is *maximal* if $T \cup p \cup \varphi(x_1, \dots, x_n)$ is unsatisfiable for any pp-formula $\varphi \notin T \cup p$. A (pp-) n -type of a structure \mathfrak{A} is just a (pp-) n -type of the theory $\text{Th}(\mathfrak{A})$. We define (maximal) ep- n -types in exactly the like fashion, with ep-formulas. A *complete* ep- n -type of a structure \mathfrak{A} is the set of all ep-formulas holding on some tuple from \mathfrak{A} . An n -type $p(x_1, \dots, x_n)$ of \mathfrak{A} is *realized* in \mathfrak{A} if there exists $a'_1, \dots, a'_n \in A$ s.t., for each $\varphi \in p$, $\mathfrak{A} \models \varphi(a'_1, \dots, a'_n)$. For an infinite cardinal κ , a structure \mathfrak{A} is κ -saturated if, for all $\beta < \kappa$ and expansions $(\mathfrak{A}; \langle a_\alpha \rangle_{\alpha < \beta})$ of \mathfrak{A} , every 1-type of $(\mathfrak{A}; \langle a_\alpha \rangle_{\alpha < \beta})$ is realized in $(\mathfrak{A}; \langle a_\alpha \rangle_{\alpha < \beta})$. We say that an infinite \mathfrak{A} is *saturated* when it is $|A|$ -saturated. Realization of pp- (ep-)types and pp- (ep-) (κ) -saturation is defined in exactly the analogous way. Note that a structure that is κ -saturated is a fortiori ep- κ -saturated. (Ep- κ -saturation and pp- κ -saturation in fact coincide, though we will only need to apply the trivial direction of this relationship.) A theory T is said to be κ -categorical, for some cardinal κ , if it has a unique model of cardinality κ , up to isomorphism. It is known that, if T is κ -categorical for one uncountable cardinal κ , then T is κ' -categorical for all uncountable cardinals κ' . A structure \mathfrak{A} , of cardinality κ , is said to be κ -categorical if $\text{Th}(\mathfrak{A})$ is κ -categorical.

Let \mathfrak{A} and \mathfrak{B} be τ -structures. A *homomorphism* from \mathfrak{A} to \mathfrak{B} is a function f from A to B such that for each k -ary relation symbol R in τ and each k -tuple (a_1, \dots, a_k) , if $(a_1, \dots, a_k) \in R^{\mathfrak{A}}$, then $(f(a_1), \dots, f(a_k)) \in R^{\mathfrak{B}}$. Also, for each constant c in τ , $f(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$. In this case we say that the map f *preserves* the relation R . Injective homomorphisms that also preserve the complement of each relation are called *embeddings*. Surjective embeddings are called *isomorphisms*; homomorphisms and isomorphisms from \mathfrak{A} to itself are called *endomorphisms* and *automorphisms*, respectively. An *elementary embedding* is an embedding that preserves all fo-formulas. A structure is a *core* if all its endomorphisms are embeddings [3] (in the finite this coincides with the more usual definition of all endomorphisms being automorphisms). \mathfrak{B} is a *core* of \mathfrak{A} if \mathfrak{B} is a core and \mathfrak{A} and \mathfrak{B} are homomorphically equivalent. We will make use later of the following lemma, a close relative of Theorem 10.3.1 in [21].

Lemma 2.1. *Let \mathfrak{A} and \mathfrak{B} be τ -structures, where \mathfrak{B} is pp- $|A|$ -saturated. Suppose f is a mapping from $\{a_\alpha : \alpha < \mu\} \subseteq A$ ($\mu < |A|$) to B such that all pp- $(\tau \cup \{c_\alpha : \alpha < \mu\})$ -sentences true on $(\mathfrak{A}; \langle a_\alpha \rangle_{\alpha < \mu})$ are true on $(\mathfrak{B}; \langle f(a_\alpha) \rangle_{\alpha < \mu})$. Then f can be extended to a homomorphism from \mathfrak{A} to \mathfrak{B} .*

Proof. Note that, if \mathfrak{B} is finite, then \mathfrak{B} is pp- $|A|$ -saturated no matter what cardinality \mathfrak{A} is.

Suppose $\mu < |A| = \kappa_A$. Let $\langle a'_\alpha \rangle_{\alpha < \kappa_A}$ well-order \mathfrak{A} such that $\langle a'_\alpha \rangle_{\alpha < \mu} = \langle a_\alpha \rangle_{\alpha < \mu}$ (there is the implicit and harmless assumption that $\langle a_\alpha \rangle_{\alpha < \mu}$ contains no repetitions). Set $\langle b'_\alpha \rangle_{\alpha < \mu} := \langle f(a_\alpha) \rangle_{\alpha < \mu}$.

We will construct by transfinite recursion on β (up to κ_A) a sequence $\langle b'_\alpha \rangle_{\alpha < \beta}$ such that we maintain the inductive hypothesis

(*) all pp- $(\tau \cup \{c_\alpha : \alpha < \beta\})$ -sentences true on $(\mathfrak{A}; \langle a'_\alpha \rangle_{\alpha < \beta})$ are true on $(\mathfrak{B}; \langle b'_\alpha \rangle_{\alpha < \beta})$.

The result will clearly then follow by reading f as the map $\{a'_\alpha \mapsto b'_\alpha\}_{\alpha < \kappa_A}$.

(Base Case.) $\beta := \mu$. Follows from hypothesis of lemma.

(Inductive Step. Limit ordinals.) $\beta := \lambda$. Property (*) holds as a sentence can only mention a finite collection of constants, whose indices must all be less than some $\gamma < \lambda$.

(Inductive Step. Successor ordinals.) $\beta := \gamma + 1 < \kappa_A$. Set

$$\Sigma := \{\varphi(x) : \varphi \text{ is a pp-}(\tau \cup \{c_\alpha : \alpha < \gamma\})\text{-formula s.t. } (\mathfrak{A}; \langle a'_\alpha \rangle_{\alpha < \gamma}) \models \varphi(a'_\gamma)\}.$$

By (*), for every $\varphi \in \Sigma$, $(\mathfrak{B}; \langle b'_\alpha \rangle_{\alpha < \gamma}) \models \exists x.\varphi(x)$. By compactness, since Σ is closed under conjunction, we have that Σ is a pp-1-type of $(\mathfrak{B}; \langle b'_\alpha \rangle_{\alpha < \gamma})$. By pp- $|A|$ -saturation of \mathfrak{B} it is realized by some element $b'_\gamma \in B$. By construction we maintain that all pp- $(\tau \cup \{c_\alpha : \alpha < \gamma + 1\})$ -sentences true on $(\mathfrak{A}; \langle a'_\alpha \rangle_{\alpha < \gamma + 1})$ are true on $(\mathfrak{B}; \langle b'_\alpha \rangle_{\alpha < \gamma + 1})$. \square

The previous lemma, as so often in this paper, holds equally for ep, in place of pp. However, we will want to apply it later to specifically the primitive positive. As a rule, we choose ep over pp in all cases where either could be used without affecting the exposition.

For a sequence of τ -structures \mathfrak{A}_α , $\alpha < \mu$, define the *direct* (or categorical) product $\prod_{\alpha < \mu} \mathfrak{A}_\alpha$ to be the τ -structure on domain $\prod_{\alpha < \mu} A_\alpha$ such that $(\langle a_\alpha^1 \rangle_{\alpha < \mu}, \dots, \langle a_\alpha^r \rangle_{\alpha < \mu}) \in R^{\prod_{\alpha < \mu} \mathfrak{A}_\alpha}$ iff $(a_\alpha^1, \dots, a_\alpha^r) \in R^{\mathfrak{A}_\alpha}$ for each $\alpha < \mu$. Also, $c^{\prod_{\alpha < \mu} \mathfrak{A}_\alpha} = \langle c^{\mathfrak{A}_\alpha} \rangle_{\alpha < \mu}$. Note that short direct products are indicated infix with \times . A property of pp-sentences φ that we will use later is that $\mathfrak{A} \models \varphi$ and $\mathfrak{B} \models \varphi$ iff $\mathfrak{A} \times \mathfrak{B} \models \varphi$.

Let $\langle \mathfrak{A} \rangle_{\text{fo}}$ (respectively, $\langle \mathfrak{A} \rangle_{\text{ep}}$ and $\langle \mathfrak{A} \rangle_{\text{pp}}$) be the sets of relations, over domain A , that are fo- (respectively, ep- and pp-) definable over \mathfrak{A} (without parameters). Let $\text{Aut}(\mathfrak{A})$ and $\text{End}(\mathfrak{A})$ be the sets of automorphisms and endomorphisms, respectively, of \mathfrak{A} . A κ -*polymorphism* of \mathfrak{A} is a homomorphism from \mathfrak{A}^κ to \mathfrak{A} , where the power is with respect to the direct product already defined. Let $\text{Pol}^\infty(\mathfrak{A})$, $\text{Pol}^\omega(\mathfrak{A})$ and $\text{Pol}(\mathfrak{A})$ be the sets of κ -polymorphisms (for any κ), κ -polymorphisms (for $\kappa \leq \omega$) and k -polymorphisms (for each finite k), respectively. For a set of operations F on domain A , define $\text{Inv}(F)$ to be the set of relations, over A , that are preserved by (invariant under) each of the operations in F (note that the condition of preservation of an m -ary relation by a κ -ary function $f : \mathfrak{A}^\kappa \rightarrow \mathfrak{A}$ is *component-wise*, i.e. if $(a_1^\beta, \dots, a_m^\beta) \in R^{\mathfrak{A}}$, for all $\beta < \kappa$, then $(f(\langle a_1^\beta \rangle_{\beta < \kappa}), \dots, f(\langle a_m^\beta \rangle_{\beta < \kappa})) \in R^{\mathfrak{A}}$).

Let \bar{t} be a k -tuple of elements from a structure \mathfrak{A} . Then the *orbit* of \bar{t} under $\text{Aut}(\mathfrak{A})$ is the set $\{f(\bar{t}) \mid f \in \text{Aut}(\mathfrak{A})\}$.

Theorem 2.2 (Ryll-Nardzewski, Engeler, Svenonius; see e.g. [21]). *A countable relational structure \mathfrak{A} is ω -categorical if and only if the automorphism group of \mathfrak{A} has for each k finitely many orbits of k -tuples.*

It is a well-known consequence of the proof of Theorem 2.2, as presented for instance in [21], that \mathfrak{A} is ω -categorical if and only if $\text{Inv}(\text{Aut}(\mathfrak{A})) = \langle \mathfrak{A} \rangle_{\text{fo}}$ (see, e.g., [2]). One direction persists in the realm of the primitive positive, as attested to by the following.

Theorem 2.3 (see [19, 12, 11]). *When \mathfrak{A} is finite or ω -categorical, $\text{Inv}(\text{Pol}(\mathfrak{A})) = \langle \mathfrak{A} \rangle_{\text{pp}}$.*

This characterization is not tight, i.e., there are infinite non- ω -categorical structures \mathfrak{A} for which $\text{Inv}(\text{Pol}(\mathfrak{A})) = \langle \mathfrak{A} \rangle_{\text{pp}}$ [29].

2.2. The constraint satisfaction problem. For a relational structure \mathfrak{A} over a finite signature, $\text{CSP}(\mathfrak{A})$ is the computational problem to decide whether a given pp-sentence is true in \mathfrak{A} . Proof of the following is straightforward (see, e.g., [23]).

Proposition 2.4. *For \mathfrak{A} and \mathfrak{A}' with the same domain, such that $\langle \mathfrak{A} \rangle_{\text{pp}} \subseteq \langle \mathfrak{A}' \rangle_{\text{pp}}$, we have that $\text{CSP}(\mathfrak{A})$ is polynomial time many-to-one reducible to $\text{CSP}(\mathfrak{A}')$.*

We remark that this result holds for logspace reductions, though this is harder to see and requires the celebrated result of [28]. In light of this observation, together with Theorem 2.3, we may use the sets $\text{Pol}(\mathfrak{A})$ to classify the computational complexity of $\text{CSP}(\mathfrak{A})$, and a most successful research program has run in this direction (see [23, 13, 14], and [17] for a survey).

Sets of the form $\text{Pol}(\mathfrak{A})$ are always *clones* (for definitions, see [35]), and the machinery of Clone Theory can be brought to bear on the classification program for CSPs (e.g., the tame congruence theory from [20] as laid out in [16]). It often transpires that instances of the CSP with low complexity can be explained by the presence of certain polymorphisms on the template. When \mathfrak{A} is finite, the class of problems $\text{CSP}(\mathfrak{A})$ is conjectured to display complexity dichotomy between those problems that are in P and those that are NP-complete (a remarkable property given the breadth of CSP problems together with the result of Ladner that NP itself does not possess the dichotomy, so long as $P \neq \text{NP}$ [25]). While the *dichotomy conjecture* was formulated independently of the algebraic method [18], a conjecture as to exactly where the boundary sits was given in the algebraic language [15].

In the case where \mathfrak{A} is infinite but ω -categorical, the connection of Theorem 2.3 has been used to good effect in the complexity classification of, e.g., temporal CSPs in [10]. In that case a dichotomy between P and NP-complete was again observed. For ω -categorical templates in general, it is known that there are structures whose CSP is undecidable [5] and of various complexities [5] (even coNP-complete). While the algebraic machinery has proved very powerful in the finite and ω -categorical cases, for infinite templates that are not ω -categorical, no such technology has thus far been developed.

3. PRIMITIVE POSITIVE DEFINABILITY OF FORMULAS

To show hardness of $\text{CSP}(\mathfrak{A})$, we often try to prove that there is a finite signature reduct \mathfrak{A}' of $\langle \mathfrak{A} \rangle_{\text{pp}}$ such that $\text{CSP}(\mathfrak{A}')$ is NP-hard. An important set of relations that contains the set of all pp-definable relations $\langle \mathfrak{A} \rangle_{\text{pp}}$ is the set of all fo-definable relations $\langle \mathfrak{A} \rangle_{\text{fo}}$. For every structure \mathfrak{A} of cardinality greater than one there are fo-definable relations yielding an NP-hard CSP, and these relations are usually good candidates for proving hardness (e.g. $R := \{0, 1\}^3 \setminus \{(0, 0, 0), (1, 1, 1)\}$ on boolean domains; $\{(a_1, a_2) : a_1 \neq a_2\}$ on finite domains A , $|A| \geq 3$; $\{(a_1, a_2, a_3) : (a_1 = a_2 \wedge a_2 \neq a_3) \vee (a_1 \neq a_2 \wedge a_2 = a_3)\}$ on infinite domains A). Therefore, it is natural and important to understand which fo-definable relations are pp-definable in \mathfrak{A} .

3.1. Equivalent templates. In this section we show that, for every problem $\text{CSP}(\mathfrak{A})$, we can find a relational structure \mathfrak{M} for which $\text{CSP}(\mathfrak{A}) = \text{CSP}(\mathfrak{M})$ where infinitary polymorphisms exactly characterize pp-definability of fo-definable relations. We will do this by building a model of $\text{Th}(\mathfrak{A})$ that is highly saturated.

Definition 3.1. A τ -structure \mathfrak{M} has the *homomorphism lifting property* if, for any finite k , $a_1, \dots, a_k \in \mathfrak{M}^\omega$ and $b_1, \dots, b_k \in \mathfrak{M}$ such that all pp- $(\tau \cup \{c_1, \dots, c_k\})$ -sentences true in $(\mathfrak{M}^\omega; a_1, \dots, a_k)$ are true in $(\mathfrak{M}; b_1, \dots, b_k)$, then there is a homomorphism $f : (\mathfrak{M}^\omega; a_1, \dots, a_k) \rightarrow (\mathfrak{M}; b_1, \dots, b_k)$.

The most natural of structures with the homomorphism lifting property are those that are of large cardinality and saturated.

Lemma 3.2. *If \mathfrak{M} is a saturated structure of cardinality $\kappa = \kappa^\omega$, then \mathfrak{M} has the homomorphism lifting property.*

Proof. This follows immediately from Lemma 2.1, since $|M^\omega| \leq |M|$. \square

We remark that the continuum has the property of Lemma 3.2 – that is $2^\omega = (2^\omega)^\omega$. On the assumption of the continuum hypothesis, we could simply work with large saturated structures; because then, given an infinite model \mathfrak{A} , we can assume the existence of an elementary extension \mathfrak{M} that is of cardinality 2^ω and saturated [22]. Without such a set-theoretic assumption, we can still construct a rather unwieldy model as follows.

Lemma 3.3. *For every τ -structure \mathfrak{A} there is an elementary extension $\mathfrak{M} \succeq \mathfrak{A}$ that is ω -saturated and has the homomorphism lifting property.*

Proof. We will build \mathfrak{M} by transfinite induction, as the union of a chain of length \aleph_1 . Set $\mathfrak{M}_0 := \mathfrak{A}$. For successor ordinals $\gamma + 1$, we take an elementary extension $\mathfrak{M}_{\gamma+1} \succeq \mathfrak{M}_\gamma$ that is $|M_\gamma|^\omega$ -saturated (such always exists, see Corollary 8.2.2 [22]). For limit ordinals λ , set $\mathfrak{M}_\lambda := \bigcup_{\alpha < \lambda} \mathfrak{M}_\alpha$; finally, let $\mathfrak{M} := \mathfrak{M}_{\aleph_1}$.

\mathfrak{M} is ω -saturated by construction. It remains to prove that \mathfrak{M} has the homomorphism lifting property. Consider the $b_1, \dots, b_k \in M$ and $a_1, \dots, a_k \in M^\omega$. The set of coordinates (of M) involved here,

$$A := \{b_1, \dots, b_k, a_1(1), a_1(2), \dots, a_2(1), a_2(2), \dots, \dots, a_k(1), a_k(2), \dots\},$$

is of size $\leq \omega$. It follows that there is some $\mu < \aleph_1$ such that $A \subseteq M_\mu$ (this is why the chain used in the construction of \mathfrak{M} is of length \aleph_1).

Suppose we are given $a_1, \dots, a_k \in \mathfrak{M}^\omega$ and $b_1, \dots, b_k \in \mathfrak{M}$ such that all pp- $(\tau \cup \{c_1, \dots, c_k\})$ -sentences true in $(\mathfrak{M}^\omega; a_1, \dots, a_k)$ are true in $(\mathfrak{M}; b_1, \dots, b_k)$. Let f_{-1} be the partial map from M^ω to M sending a_1, \dots, a_k to b_1, \dots, b_k (well-defined as all equalities are pp). We first argue that all pp- $(\tau \cup \{c_1, \dots, c_k\})$ -sentences true in $(\mathfrak{M}_\mu^\omega; a_1, \dots, a_k)$ are true in $(\mathfrak{M}_{\mu+1}; f_{-1}(a_1), \dots, f_{-1}(a_k))$. Let φ be such a sentence. Then:

$$\begin{array}{llll} & (\mathfrak{M}_\mu^\omega; a_1, \dots, a_k) \models \varphi & \Rightarrow & \\ \text{for each } i & (\mathfrak{M}_\mu; a_1(i), \dots, a_k(i)) \models \varphi & \Rightarrow^a & \\ \text{for each } i & (\mathfrak{M}; a_1(i), \dots, a_k(i)) \models \varphi & \Rightarrow & \\ & (\mathfrak{M}^\omega; a_1, \dots, a_k) \models \varphi & \Rightarrow^b & \\ & (\mathfrak{M}; f_{-1}(a_1), \dots, f_{-1}(a_k)) \models \varphi & \Rightarrow^c & \\ & (\mathfrak{M}_{\mu+1}; f_{-1}(a_1), \dots, f_{-1}(a_k)) \models \varphi. & & \end{array}$$

(a) since $\mathfrak{M}_\mu \preceq \mathfrak{M}$; (b) by hypothesis; (c) since $\mathfrak{M} \succeq \mathfrak{M}_{\mu+1}$. It follows from Lemma 2.1 that there is a homomorphism $f_0 : \mathfrak{M}_\mu^\omega \rightarrow \mathfrak{M}_{\mu+1}$ extending f_{-1} .

We will now proceed with a transfinite induction up to \aleph_1 . For successor ordinals, $\gamma + 1$, suppose that we have a homomorphism $f_\gamma : \mathfrak{M}_{\mu+\gamma}^\omega \rightarrow \mathfrak{M}_{\mu+\gamma+1}$. We will build a homomorphism $f_{\gamma+1} : \mathfrak{M}_{\mu+\gamma+1}^\omega \rightarrow \mathfrak{M}_{\mu+\gamma+2}$ extending f_γ . Since f_γ is a homomorphism, all pp- $(\tau \cup \{c_\alpha : \alpha < |M_{\mu+\gamma}|^\omega\})$ -sentences true in $(\mathfrak{M}_{\mu+\gamma}^\omega; \langle a_\alpha \rangle_{\alpha < |M_{\mu+\gamma}|^\omega})$ – where $\langle a_\alpha \rangle_{\alpha < |M_{\mu+\gamma}|^\omega}$ well-orders the elements of $\mathfrak{M}_{\mu+\gamma}^\omega$ – are true in $(\mathfrak{M}_{\mu+\gamma+1}; \langle f_\gamma(a_\alpha) \rangle_{\alpha < |M_{\mu+\gamma}|^\omega})$. It follows that all pp- $(\tau \cup \{c_\alpha : \alpha < |M_{\mu+\gamma}|^\omega\})$ -sentences true in $(\mathfrak{M}_{\mu+\gamma+1}^\omega; \langle a_\alpha \rangle_{\alpha < |M_{\mu+\gamma}|^\omega})$ are true in $(\mathfrak{M}_{\mu+\gamma+2}; \langle f(a_\alpha) \rangle_{\alpha < |M_{\mu+\gamma}|^\omega})$ – let φ be such a sentence, we give the derivation again:

$$\begin{array}{llll} & (\mathfrak{M}_{\mu+\gamma+1}^\omega; \langle a_\alpha \rangle_{\alpha < |M_{\mu+\gamma}|^\omega}) \models \varphi & \Rightarrow & \\ \text{for each } i & (\mathfrak{M}_{\mu+\gamma+1}; \langle a_\alpha(i) \rangle_{\alpha < |M_{\mu+\gamma}|^\omega}) \models \varphi & \Rightarrow & \\ \text{for each } i & (\mathfrak{M}_{\mu+\gamma}; \langle a_\alpha(i) \rangle_{\alpha < |M_{\mu+\gamma}|^\omega}) \models \varphi & \Rightarrow & \\ & (\mathfrak{M}_{\mu+\gamma}^\omega; \langle a_\alpha \rangle_{\alpha < |M_{\mu+\gamma}|^\omega}) \models \varphi & \Rightarrow & \\ & (\mathfrak{M}_{\mu+\gamma+1}; \langle f(a_\alpha) \rangle_{\alpha < |M_{\mu+\gamma}|^\omega}) \models \varphi & \Rightarrow & \\ & (\mathfrak{M}_{\mu+\gamma+2}; \langle f(a_\alpha) \rangle_{\alpha < |M_{\mu+\gamma}|^\omega}) \models \varphi. & & \end{array}$$

Now we can use Lemma 2.1 to derive some homomorphism $f_{\gamma+1} : \mathfrak{M}_{\mu+\gamma+1}^\omega \rightarrow \mathfrak{M}_{\mu+\gamma+2}$ extending f_γ . For limit ordinals λ , set $f_\lambda := \bigcup_{\alpha < \lambda} f_\alpha$.

Finally, we arrive at the homomorphism $f_{\aleph_1} : \mathfrak{M}^\omega \rightarrow \mathfrak{M}$, which has the desired property. \square

We describe the model built as in the previous lemma as a *monster* model (this term is used with different connotation in set theory and model theory).

Let $\langle \mathfrak{A} \rangle_{\text{pp}\infty}$ be the set of relations pp-definable on \mathfrak{A} , *possibly involving infinitary conjunction* (of pp-formulas in a finite number of free variables). Because we will use it again later, we give the following lemma in its strongest form.

Lemma 3.4. *For all structures \mathfrak{A} , $\langle \mathfrak{A} \rangle_{\text{pp}\infty} \subseteq \text{Inv}(\text{Pol}^\infty(\mathfrak{A}))$.*

Proof. We argue by induction on the term-complexity of the formula. Let $f : A^\alpha \rightarrow A$ be a polymorphism of \mathfrak{A} .

(Base Case.) $\varphi(\bar{v}) := R(\bar{v})$. Trivial.

(Inductive Step.) There are two subcases. In the following, suppose \bar{v} is an m -tuple. Let $\langle \bar{a}_\beta \rangle_{\beta < \alpha}$, be a sequence of m -tuples from \mathfrak{A} such that $\mathfrak{A} \models \varphi(\bar{a}_\beta)$, for all β .

(Existential Quantification.) $\varphi(\bar{v}) := \exists u. \psi(\bar{v}, u)$. Suppose we have $\mathfrak{A} \models \varphi(\bar{a}_\beta)$ for each $\beta < \alpha$. From each $\mathfrak{A} \models \exists u. \psi(\bar{a}_\beta, u)$, derive the witness a'_β for u and use the inductive hypothesis to deduce that $\mathfrak{A} \models \psi(f(\langle \bar{a}_\beta \rangle_{\beta < \alpha}), f(\langle a'_\beta \rangle_{\beta < \alpha}))$. It follows that $\mathfrak{A} \models \exists u. \psi(f(\langle \bar{a}_\beta \rangle_{\beta < \alpha}), u)$ and we are able to deduce $\mathfrak{A} \models \varphi(f(\langle \bar{a}_\beta \rangle_{\beta < \alpha}))$.

(Infinite Conjunctions.) $\varphi(\bar{v}) := \bigwedge_{\mu < \gamma} \psi_\mu(\bar{v})$. Suppose we have $\mathfrak{A} \models \varphi(\bar{a}_\beta)$ for each $\beta < \alpha$. Then for each $\mu < \gamma$ and $\beta < \alpha$ we have $\mathfrak{A} \models \psi_\mu(\bar{a}_\beta)$. By inductive hypothesis, we have each $\mathfrak{A} \models \psi_\mu(f(\langle \bar{a}_\beta \rangle_{\beta < \alpha}))$. The result $\mathfrak{A} \models \varphi(f(\langle \bar{a}_\beta \rangle_{\beta < \alpha}))$ follows. \square

We are now ready for the main result of this section.

Theorem 3.5. *Let \mathfrak{A} be a structure, over a countable¹ signature, that is ω -saturated and has the homomorphism lifting property. Then a fo-definable relation R is preserved by the ω -polymorphisms of \mathfrak{A} if and only if R is pp-definable in \mathfrak{A} , i.e.*

$$\text{Inv}(\text{Pol}^\omega(\mathfrak{A})) \cap \langle \mathfrak{A} \rangle_{\text{fo}} = \langle \mathfrak{A} \rangle_{\text{pp}}.$$

Proof. (Backwards.) That pp-formulas are preserved by ω -polymorphisms in any structure is a special case of Lemma 3.4.

(Forwards.) Suppose that R is a k -ary relation that is preserved by all ω -polymorphisms of \mathfrak{A} and that has a first-order definition φ in \mathfrak{A} . Let

$$\Psi := \{\psi(x_1, \dots, x_k) : \psi \text{ is a pp-}\tau\text{-formula s.t. } \mathfrak{A} \models \varphi(x_1, \dots, x_k) \rightarrow \psi(x_1, \dots, x_k)\}.$$

We first show, for all $b_1, \dots, b_k \in A$, that, if $\mathfrak{A} \models \Psi(b_1, \dots, b_k)$, then $\mathfrak{A} \models \varphi(b_1, \dots, b_k)$.

Take $b_1, \dots, b_k \in A$ s.t. $\mathfrak{A} \models \psi(b_1, \dots, b_k)$ for each $\psi \in \Psi$; if such elements do not exist there is nothing to show. Let U be the set of all pp- τ -formulas $\theta(x_1, \dots, x_k)$ such that $\mathfrak{A} \models \neg\theta(b_1, \dots, b_k)$. U contains \perp , $\perp \wedge \perp$ etc., and so we may assume U to be countably infinite. Let $(\theta_i)_{i < \omega}$ be an enumeration of U . We claim that for every $\theta_i \in U$ there exists a k -tuple $\bar{a}^{\theta_i} := (a_1^{\theta_i}, \dots, a_k^{\theta_i})$ from A such that $\mathfrak{A} \models \neg\theta_i(a_1^{\theta_i}, \dots, a_k^{\theta_i}) \wedge \varphi(a_1^{\theta_i}, \dots, a_k^{\theta_i})$. Otherwise, $\mathfrak{A} \models \varphi(x_1, \dots, x_k) \rightarrow \theta_i(x_1, \dots, x_k)$, and we derive $\theta_i \in \Psi$ and the consequent contradiction $\mathfrak{A} \models \theta_i(b_1, \dots, b_k)$.

¹ The restriction to countable signatures allows us to consider polymorphisms of at most arity ω . Were we to consider signatures of infinite cardinality α then the theorem would hold with polymorphisms of arity α (under a slightly different definition of the homomorphism lifting property involving α in the exponent).

Consider the k -tuple $\bar{a} := \prod_{i < \omega} \bar{a}^{\theta_i}$ in \mathfrak{A}^ω . Observe that every pp- τ -formula $\chi(x_1, \dots, x_k)$ s.t. $\mathfrak{A}^\omega \models \chi(\bar{a})$ is s.t. $\mathfrak{A} \models \chi(b_1, \dots, b_k)$. To see this, suppose that $\mathfrak{A} \models \neg\chi(b_1, \dots, b_k)$. Therefore $\chi \in U$, and by choice of \bar{a}^χ we have $\mathfrak{A} \models \neg\chi(\bar{a}^\chi)$. But then $\mathfrak{A}^\omega \models \neg\chi(\bar{a})$.

Now, we have just shown that all pp- $(\tau \cup \{c_1, \dots, c_k\})$ -sentences that hold on $(\mathfrak{A}^\omega; \bar{a})$ also hold on $(\mathfrak{A}; b_1, \dots, b_k)$. Since \mathfrak{A} has the homomorphism lifting property, the existence of a homomorphism $f : (\mathfrak{A}^\omega; \bar{a}) \rightarrow (\mathfrak{A}; b_1, \dots, b_k)$ follows from our definitions. But f is an ω -polymorphism of \mathfrak{A} , which preserves φ , and hence we derive $\mathfrak{A} \models \varphi(b_1, \dots, b_k)$.

It remains to be shown that Ψ is equivalent on \mathfrak{A} to a single pp-formula. Note that $\Psi(c_1, \dots, c_k) \cup \{\neg\varphi(c_1, \dots, c_k)\} \cup \text{Th}(\mathfrak{A})$ is unsatisfiable; for otherwise there is a $\mathfrak{B} \models \text{Th}(\mathfrak{A})$ and $b'_1, \dots, b'_k \in B$, s.t. $(\mathfrak{B}; b'_1, \dots, b'_k) \models \Psi(c_1, \dots, c_k)$ and $(\mathfrak{B}; b'_1, \dots, b'_k) \models \neg\varphi(c_1, \dots, c_k)$. Since \mathfrak{A} is ω -saturated, this yields some $b''_1, \dots, b''_k \in A$ such that both $(\mathfrak{A}; b''_1, \dots, b''_k) \models \Psi(c_1, \dots, c_k)$ and $(\mathfrak{A}; b''_1, \dots, b''_k) \models \neg\varphi(c_1, \dots, c_k)$, which is a contradiction. By compactness of first-order logic there is a finite subset Ψ' of Ψ such that $\Psi'(c_1, \dots, c_k) \cup \{\neg\varphi(c_1, \dots, c_k)\} \cup \text{Th}(\mathfrak{A})$ is unsatisfiable, i.e. $\Psi'(c_1, \dots, c_k) \cup \text{Th}(\mathfrak{A}) \models \varphi(c_1, \dots, c_k)$. Set $\psi'(x_1, \dots, x_k) := \bigwedge_{\psi \in \Psi'} \psi(x_1, \dots, x_k)$, to derive $\text{Th}(\mathfrak{A}) \models \psi'(x_1, \dots, x_k) \rightarrow \varphi(x_1, \dots, x_k)$. Since $\mathfrak{A} \models \varphi(x_1, \dots, x_k) \rightarrow \psi'(x_1, \dots, x_k)$ by construction, the result follows. \square

Corollary 3.6. *Let \mathfrak{A} be any structure with finite relational signature. Then there exists a structure \mathfrak{M} such that $\text{CSP}(\mathfrak{A}) = \text{CSP}(\mathfrak{M})$, and such that an fo-definable relation R is pp-definable in \mathfrak{M} if and only if R is preserved by all ω -polymorphisms of \mathfrak{M} .*

Proof. By Lemma 3.3, there is a monster elementary extension of $\mathfrak{M} \succeq \mathfrak{A}$ with the homomorphism lifting property. We now apply Theorem 3.5. \square

In the parlance of [27], the following may be seen as the “global” analog of Theorem 3.5.

Corollary 3.7. *An fo-formula φ is preserved by the ω -polymorphisms of all elementary extensions of \mathfrak{A} if and only if φ is pp-definable in \mathfrak{A} .*

Proof. (Backwards.) Follows from Lemma 3.4.

(Forwards.) Since φ is preserved by the ω -polymorphisms of the monster elementary extension $\mathfrak{M} \succeq \mathfrak{A}$ constructed in Lemma 3.3, it follows from Theorem 3.5 that φ is pp-definable on \mathfrak{M} . But this is a fortiori a pp-definition on \mathfrak{A} . \square

Corollary 3.8. *Let T be an uncountably categorical fo-theory, and \mathfrak{A} a model of T of cardinality $\kappa = \kappa^\omega$. Then $\text{Inv}(\text{Pol}^\omega(\mathfrak{A})) \cap \langle \mathfrak{A} \rangle_{\text{fo}} = \langle \mathfrak{A} \rangle_{\text{pp}}$.*

Proof. It is well-known that uncountable models of uncountably categorical theories are saturated in their own cardinality (Fact 1.2. in [36]). Hence, the statement follows from Lemma 3.2 and Theorem 3.5. \square

3.2. Tightness of Theorem 3.5. One might be interested in the following potential strengthenings of Theorem 3.5.

- (1) To derive the statement for arbitrary relations (not just for fo-definable relations).
- (2) To assume preservation under finitary polymorphisms (not infinitary polymorphisms).
- (3) To show the statement for arbitrary models of $\text{Th}(\mathfrak{A})$ (not just for structures with the homomorphism lifting property).

The following proposition shows that each of these assumptions is necessary.

Proposition 3.9.

- (1) There is a saturated structure $\mathfrak{A}_{\text{sat}}$ of cardinality 2^ω such that $\text{Inv}(\text{Pol}^\omega(\mathfrak{A}_{\text{sat}})) \neq \langle \mathfrak{A}_{\text{sat}} \rangle_{\text{pp}}$.
- (2) There is a saturated structure $\mathfrak{A}_{\text{sat}}$ of cardinality 2^ω such that $\text{Inv}(\text{Pol}(\mathfrak{A}_{\text{sat}})) \cap \langle \mathfrak{A}_{\text{sat}} \rangle_{\text{fo}} \neq \langle \mathfrak{A}_{\text{sat}} \rangle_{\text{pp}}$.
- (3) There is a structure \mathfrak{A} such that $\text{Inv}(\text{Pol}^\omega(\mathfrak{A})) \cap \langle \mathfrak{A} \rangle_{\text{fo}} \neq \langle \mathfrak{A} \rangle_{\text{pp}}$.

Proof. Necessity of intersection with fo. Let us consider the model $\mathfrak{A} := (\mathbb{Q}; +, 1, (u = v \vee x = y))$. By Lemma 5.3, the infinitary polymorphisms of this structure are equivalent to its endomorphisms, and, in the presence of a fixed 1, it can easily be seen that its only endomorphism is the identity (indeed, there is a pp-definition of each of the rationals from 1 and +). It follows that all subsets of \mathbb{Q} are in $\text{Inv}(\text{Pol}^\omega(\mathfrak{A}))$, yet $\langle \mathfrak{A} \rangle_{\text{pp}}$ must be countable. Of course, \mathfrak{A} is neither saturated nor of cardinality 2^ω . But the continuum of subsets of \mathbb{Q} will remain Inv-Pol^ω in a saturated model of $\text{Th}(\mathfrak{A})$ of such cardinality (a copy of \mathfrak{A} sits elementary in all models of its theory). The existence of a saturated model of $\text{Th}(\mathfrak{A})$ of cardinality 2^ω follows from this theory's strong minimality (Fact 1.2. in [36]).

Necessity of infinitary polymorphisms. Let $\{U_i : i \in \omega\}$ be a set of unary relations. Consider the model $\mathfrak{A} := (\mathbb{N}; U_i : i \in \omega)$, involving a countable set of unary relations, defined by $U_i := \mathbb{N} \setminus \{0, i\}$. Diagrammatically,

	U_1	U_2	U_3	\dots
0	×	×	×	\dots
1	×	√	√	\dots
2	√	×	√	\dots
3	√	√	×	\dots
⋮	⋮	⋮	⋮	

Consider the fo-definable unary relation $P(v) := U_1(v) \vee U_2(v)$, i.e. $P := \mathbb{N} \setminus \{0\}$. It is straightforward to verify that P is closed under the finitary polymorphisms of \mathfrak{A} and is not pp-definable over \mathfrak{A} . Note that P is not preserved under the infinitary polymorphism $f : \mathbb{N}^\omega \rightarrow \mathbb{N}$ of \mathfrak{A} defined by $f(\bar{w}) = 0$, if \bar{w} contains all elements of $\mathbb{N} \setminus \{0\}$, and $f(\bar{w}) = w_0$ (the first element of the sequence \bar{w}), otherwise. Again, these properties will remain if we move to a saturated model $\mathfrak{A}_{\text{sat}}$ of cardinality 2^ω (such a model will simply be \mathfrak{A} augmented with a continuum of elements for which all of the relations $\{U_i : i \in \omega\}$ hold).

We now detail a finite signature variant of the above structure that also serves as a suitable (counter)example. Consider the signature $\langle E, R \rangle$ involving two binary relations, edge and red edge. Let the structure \mathfrak{A} contain

- a directed ω - E -path: i.e., vertices $\{(0, i) : i < \omega\}$ and E -edges $\{((0, i), (0, i + 1)) : i < \omega\}$, and for each $j < \omega$:
- a directed ω - E -path with overlaid undirected R -path omitting only the j th edge: i.e. vertices $\{(j, i) : i < \omega\}$ with E -edges $\{((j, i), (j, i + 1)) : i < \omega\}$ and R -edges $\{((j, i), (j, i + 1)), ((j, i + 1), (j, i)) : i < \omega, i + 1 \neq j\}$.

Consider the first-order definable unary relation $P(v) := \exists x, y. R(v, x) \vee (E(v, x) \wedge R(x, y))$. It is not hard to verify that P is preserved by the finitary polymorphisms of \mathfrak{A} , but is not pp-definable over \mathfrak{A} (as it is not preserved by the ω -polymorphisms of \mathfrak{A}). The saturated elementary extension of \mathfrak{A} of cardinality 2^ω contains 2^ω copies of each one of the countably infinite set of paths just described. It remains easy to see that P is preserved by $\mathfrak{A}_{\text{sat}}$'s finitary but not ω -ary polymorphisms.

Necessity of highly saturated structures. Consider the structure $\mathfrak{A} := (\mathbb{Q}; x = 1, x < 0, S_2(x, y))$, where $S_2 := \{(x, y) : 2x < y, 0 < y \leq 1\}$. Now, $x \leq 0$ is clearly fo-definable in \mathfrak{A} (this is from $x < 0$; actually $x \leq 0$ is fo-definable from just S_2 , so $x < 0$ is not needed as an extensional relation). It is also in $\text{Inv}(\text{Pol}^\omega(\mathfrak{A}))$, being definable by the following infinite conjunction of pp-formulas in one free variable (see Lemma 3.4).

$$\bigwedge_{i \in \omega} \exists z \exists y_1, \dots, y_i. S_2(x, y_1) \wedge S_2(y_1, y_2) \wedge \dots \wedge S_2(y_i, z) \wedge z = 1.$$

We will now argue that it is not pp-definable.

Lemma. Let $\bar{x} := (x_1, \dots, x_k)$ and suppose that $\varphi(\bar{x}) \in \langle \mathfrak{A} \rangle_{\text{pp}}$. If $\mathfrak{A} \models \varphi(\bar{a})$ and $a_{\lambda_1}, \dots, a_{\lambda_j}$ list exactly the elements of \bar{a} that are 0, then there exists $\epsilon > 0$ such that, for all $\epsilon \geq \delta \geq 0$, $\mathfrak{A} \models \varphi(\bar{a}[a_{\lambda_1}/\delta, \dots, a_{\lambda_j}/\delta])$.

Proof. By induction on the term complexity of φ .

(Base Cases.) φ is an atom. The statement is trivially true if $\varphi(x) := x = 1, x < 0$ or $x = x$. Suppose $\varphi(x_1, x_2) := S_2(x_1, x_2)$; if $S_2(a_1, a_2)$, then only a_1 could be zero. Set $\epsilon := a_2/2$.

(Inductive Step.) There are two subcases.

$\varphi(\bar{x}) := \psi_1(\bar{x}) \wedge \psi_2(\bar{x})$. There exist respective witnesses ϵ_1 and ϵ_2 for $\psi_1(\bar{a})$ and $\psi_2(\bar{a})$; we may set $\epsilon := \min\{\epsilon_1, \epsilon_2\}$ as the witness for $\varphi(\bar{a})$.

$\varphi(\bar{x}) := \exists y. \psi(y, \bar{x})$. If $\mathfrak{A} \models \varphi(\bar{a})$ holds, then we may choose a b s.t. $\mathfrak{A} \models \psi(b, \bar{a})$. By inductive hypothesis, there exists an appropriate ϵ for $\psi(b, \bar{a})$ and this may also be used for $\varphi(\bar{a})$. \square

That $x \leq 0$ is not pp-definable is a trivial consequence of the lemma, for suppose it were defined by $\varphi(x)$. Since $\varphi(0)$ holds, we may derive the contradiction that $\varphi(\epsilon)$ holds for some $\epsilon > 0$. Note that the first part of the inductive step in the previous lemma would fail for infinite conjunctions. Finally, suppose $\mathfrak{A}_{\text{sat}}$ were a saturated model of $\text{Th}(\mathfrak{A})$ of cardinality 2^ω . While we have $\langle \mathfrak{A} \rangle_{\text{fo}} \cap \text{Inv}(\text{Pol}^\omega(\mathfrak{A})) \neq \langle \mathfrak{A} \rangle_{\text{pp}}$, we must have $\langle \mathfrak{A}_{\text{sat}} \rangle_{\text{fo}} \cap \text{Inv}(\text{Pol}^\omega(\mathfrak{A}_{\text{sat}})) = \langle \mathfrak{A}_{\text{sat}} \rangle_{\text{pp}}$. We note that $x \leq 0$ is not $\text{Inv}(\text{Pol}^\omega(\mathfrak{A}_{\text{sat}}))$. \square

4. EXISTENTIAL-POSITIVELY CLOSED MODELS

In this section, we define existential-positively closed models and demonstrate their applications in constraint satisfaction.

4.1. Definitions and basic results. We begin by stating some basic facts regarding existential-positively closed models. They are the positive analogs of existentially closed models (the latter are treated in great detail in Section 7 of [22], which is Section 8 of [21]), and have been studied under the name of existentially closed models in a recent paper on positive model theory by Ben-Yaacov [1].

Definition 4.1. A model \mathfrak{A} is *existential-positively closed for T* – or short *epc* – iff $\mathfrak{A} \models T$ and for any homomorphism h from \mathfrak{A} into another model \mathfrak{B} of T , any tuple \bar{a} from A , and any existential positive formula φ with $\mathfrak{B} \models \varphi(h(\bar{a}))$ we have that $\mathfrak{A} \models \varphi(\bar{a})$.

Note that we could equivalently have used primitive positive formulas in the previous definition. To show the existence of certain epc models we apply the direct limit construction.

Definition 4.2. Let τ be a relational signature, and let $(\mathfrak{A}_i)_{i < \kappa}$ be a sequence of τ -structures of cardinality $\leq \kappa$ such that there are homomorphisms $f_{ij} : \mathfrak{A}_i \rightarrow \mathfrak{A}_j$ with $f_{jk} \circ f_{ij} = f_{ik}$ for every $i \leq j \leq k$. Then the *direct limit* $\lim_{i < \kappa} \mathfrak{A}_i$ (with respect to the homomorphisms f_{ij}) is the τ -structure \mathfrak{A} defined as follows. The domain A of \mathfrak{A} comprises the equivalence classes of the equivalence relation \sim defined on $\bigcup_{i < \kappa} A_i$ by setting $x_i \sim x_j$ for $x_i \in A_i$ and $x_j \in A_j$ iff there is a k such that $f_{ik}(x_i) = f_{jk}(x_j)$. Let $g_i : A_i \rightarrow A$ be the *limit homomorphism*, i.e. the function that maps $a \in A_i$ to the equivalence class of a in A . For $R \in \tau$, define $\mathfrak{A} \models R(\bar{a})$ iff there is a k and $\bar{b} \in A_k$ such that $\mathfrak{A}_k \models R(\bar{b})$ and $\bar{a} = g_k(\bar{b})$. Note that limit homomorphisms behave well with respect to their homomorphism family (i.e. $g_i = g_j \circ f_{ij}$).

The direct limits defined above can be seen as a positive variant of the basic model-theoretic notion of a *union of chains* (see Section 2.4 in [21]); we essentially replace embeddings in chains by homomorphisms. Unions of chains preserve $\forall\exists$ -sentences; the analogous statement for direct limits is as follows. A sentence is called $\forall\exists^+$ if it is a universally quantified positive boolean combination of existential positive formulas and negated atomic formulas. We could have equivalently defined $\forall\exists^+$ -formulas as conjunctions of universally quantified disjunctions of primitive positive formulas and negated atomic formulas. It is easy to see that every $\forall\exists^+$ -formula can be re-written into such a formula.

Proposition 4.3 (see Theorem 2.4.6 in [21]). *Let \mathfrak{A} be the direct limit of $(\mathfrak{A}_i)_{i < \kappa}$; if φ is $\forall\exists^+$ such that $\mathfrak{A}_i \models \varphi$ for all i , then $\mathfrak{A} \models \varphi$.*

Lemma 4.4. *The class of all epc models for a theory T is closed under direct limits.*

Proof. Suppose that $\mathfrak{A} = \lim_{\alpha < \kappa} \mathfrak{A}_\alpha$ for a sequence $(\mathfrak{A}_\alpha)_{\alpha < \kappa}$ of models that are epc for T , \bar{a} is a tuple from \mathfrak{A} , φ an ep-formula, and h is a homomorphism from \mathfrak{A} into another model of T such that $\mathfrak{B} \models \varphi(h(\bar{a}))$. Then there exists a $\alpha < \kappa$ such that $\bar{a} = g_\alpha(\bar{a}')$ for \bar{a}' from \mathfrak{A}_α (where g_α is as in the definition of direct limits). Note that $h \circ g_\alpha$ is a homomorphism from \mathfrak{A}_α to \mathfrak{B} , and since \mathfrak{A}_α is a model epc for T , $\mathfrak{A}_\alpha \models \varphi(\bar{a}')$. Since g_α preserves ep-formulas, we thus also have that $\mathfrak{A} \models \varphi(\bar{a})$. \square

Proposition 4.5 (Essentially from [1]). *Let \mathfrak{A} be a model of cardinality κ of a set T of $\forall\exists^+$ sentences. Then there is a homomorphism from \mathfrak{A} to a model \mathfrak{B} that is epc for T of cardinality $\leq \kappa$.*

Proof. Set $\mathfrak{B}_0 := \mathfrak{A}$. Let $i > 0$ be a natural number and assume we have already constructed \mathfrak{B}_{i-1} of cardinality $\leq \kappa$. Let $\{(\varphi_\alpha, \bar{a}_\alpha) \mid \alpha < \kappa\}$ be a (not necessarily injective) enumeration of all pairs (φ, \bar{a}) where φ is existential positive with free variables x_1, \dots, x_n , and \bar{a} is an n -tuple from B_{i-1} . We construct a sequence $(\mathfrak{B}_i^\alpha)_{0 \leq \alpha < \kappa}$ of models of T of cardinality $\leq \kappa$ and a coherent sequence $(f_i^{\mu, \alpha})_{0 \leq \mu < \alpha < \kappa}$ of homomorphisms, where $f_i^{\mu, \alpha} : \mathfrak{B}_i^\mu \rightarrow \mathfrak{B}_i^\alpha$, as follows (coherent in the sense that, for $\alpha < \beta < \gamma$, $f_i^{\alpha, \gamma} = f_i^{\beta, \gamma} \circ f_i^{\alpha, \beta}$).

Set $\mathfrak{B}_i^0 = \mathfrak{B}_{i-1}$. Now let $\alpha = \beta + 1 < \kappa$ be a successor ordinal. Let \bar{b}_β be the image of \bar{a}_β in \mathfrak{B}_i^β under $f_i^{0, \beta}$ (in case $\beta = 0$, set $\bar{b}_0 := \bar{a}_0$). If there is a model \mathfrak{C} of T and a homomorphism $h : \mathfrak{B}_i^\beta \rightarrow \mathfrak{C}$ such that $\mathfrak{C} \models \varphi_\beta(h(\bar{b}_\beta))$, then by the theorem of Löwenheim-Skolem there is also a model \mathfrak{C}' of cardinality $\leq \kappa$ of T and a homomorphism $h' : \mathfrak{B}_i^\beta \rightarrow \mathfrak{C}'$ such that $\mathfrak{C}' \models \varphi_\beta(h'(\bar{b}_\beta))$. Set $\mathfrak{B}_i^\alpha = \mathfrak{C}'$ and $f_i^{\mu, \alpha} = h' \circ f_i^{\mu, \beta}$. Otherwise, if there is no such model \mathfrak{C} , we set $\mathfrak{B}_i^\alpha = \mathfrak{B}_i^\beta$ and $f_i^{\beta, \alpha} = \text{id}$ (the identity mapping) and $f_i^{\mu, \alpha} = f_i^{\mu, \beta}$. Finally, for limit ordinals $\alpha < \kappa$, set $\mathfrak{B}_i^\alpha = \lim_{\mu < \alpha} \mathfrak{B}_i^\mu$ and let $f_i^{\mu, \alpha}$ be the corresponding limit homomorphism from \mathfrak{B}_i^μ to \mathfrak{B}_i^α .

Let \mathfrak{B}_i be $\lim_{\alpha < \kappa} \mathfrak{B}_i^\alpha$ (with respect to $f_i^{\mu, \alpha}$) and let $g_{i-1} : \mathfrak{B}_{i-1} = \mathfrak{B}_i^0 \rightarrow \mathfrak{B}_i$ be the corresponding limit homomorphism.

Let $\mathfrak{B} = \lim_{i < \omega} \mathfrak{B}_i$. By construction, \mathfrak{B} is of cardinality $\leq \kappa$, and it is a model of T by Proposition 4.3; let $h_i : \mathfrak{B}_i \rightarrow \mathfrak{B}$ for $i < \omega$ be the corresponding homomorphisms.

The structure \mathfrak{B} is epc for T . To verify this, let g be a homomorphism from \mathfrak{B} to a model \mathfrak{C} of T , and suppose that there is a tuple \bar{b} over B and an existential positive formula φ such that $\mathfrak{C} \models \varphi(g(\bar{b}))$. There is an $i < \omega$ and an $\bar{a} \in B_i$ such that $h_i(\bar{a}) = \bar{b}$. Then $g \circ h_i$ is a homomorphism from \mathfrak{B}_i to \mathfrak{C} , and by our construction we have that $\mathfrak{B}_{i+1} \models \varphi(g_i(\bar{a}))$. Note that $h_{i+1} \circ g_i = h_i$. Thus, since h_{i+1} preserves existential positive formulas, we also have that $\mathfrak{B} \models \varphi(\bar{b})$, which is what we had to show. \square

Proposition 4.6. *Let T be a theory, and let \mathfrak{A} be a model of T . Then \mathfrak{A} is epc for T if and only if every complete ep- n -type of \mathfrak{A} is a maximal ep-type of T .*

Proof. (Forwards.) Suppose $p(x_1, \dots, x_n)$ is the complete ep- n -type, realized in \mathfrak{A} by the tuple (a_1, \dots, a_n) . Let c_1, \dots, c_n be new constant symbols that denote a_1, \dots, a_n in \mathfrak{A} . Let $\varphi(x_1, \dots, x_n)$ be an existential positive formula such that $T \cup p(c_1, \dots, c_n) \cup \{\varphi(c_1, \dots, c_n)\}$ has a model $(\mathfrak{C}; b_1, \dots, b_n)$. Now, let $(\mathfrak{C}_{\text{sat}}; b_1, \dots, b_n)$ be an $|A|$ -saturated model of $\text{Th}(\mathfrak{C}; b_1, \dots, b_n)$; such a model always exists (see Corollary 8.2.2 [22]). Clearly $(\mathfrak{C}_{\text{sat}}; b_1, \dots, b_n)$ is ep- $|A|$ -saturated, and all existential positive formulas that are true on $(\mathfrak{A}; a_1, \dots, a_n)$ are true on $(\mathfrak{C}_{\text{sat}}; b_1, \dots, b_n)$. By Lemma 2.1, there is a homomorphism h from $(\mathfrak{A}; a_1, \dots, a_n)$ to $(\mathfrak{C}_{\text{sat}}; b_1, \dots, b_n)$. Now, since $\varphi(b_1, \dots, b_n)$ holds on $\mathfrak{C}_{\text{sat}}$ and \mathfrak{A} is epc for T , we find that $\mathfrak{A} \models \varphi(a_1, \dots, a_n)$, and conclude that p is a maximal ep-type of T .

(Backwards.) Take $\mathfrak{B} \models T$, $h : \mathfrak{A} \rightarrow \mathfrak{B}$ a homomorphism, \bar{a} a tuple of elements of \mathfrak{A} , and $\varphi(x_1, \dots, x_n)$ an existential positive formula such that $\mathfrak{B} \models \varphi(h(\bar{a}))$. Let p be the ep-type of \bar{a} in \mathfrak{A} . Since \mathfrak{B} is a model of T and h preserves all ep-formulas, it follows that $T \cup p \cup \{\varphi\}$ is satisfiable. By maximality of p , we have that $\varphi \in p$, and therefore $\mathfrak{A} \models \varphi(\bar{a})$. \square

We conclude this section by noting that epc structures are related to the concept of cores, which play such an important role in the classification program for CSPs when the template is finite or ω -categorical.

Proposition 4.7. *If \mathfrak{A} is epc for $\text{Th}_{\forall^-}(\mathfrak{A})$, then \mathfrak{A} is a core. If \mathfrak{A} is ep-saturated (or finite) and expanded by all ep-definable relations, then the converse holds also.*

Proof. Suppose \mathfrak{A} is epc for $\text{Th}_{\forall^-}(\mathfrak{A})$. Take a homomorphism $h : \mathfrak{A} \rightarrow \mathfrak{A}$. By epc, for a_1, \dots, a_k in A , if $\mathfrak{A} \models R(h(a_1), \dots, h(a_k))$ or $\mathfrak{A} \models h(a_1) = h(a_2)$, then $\mathfrak{A} \models R(a_1, \dots, a_k)$ or $\mathfrak{A} \models a_1 = a_2$, respectively. It follows that h is an embedding.

Now suppose that \mathfrak{A} is ep-saturated (or finite) and expanded by all ep-definable relations. Suppose $\mathfrak{B} \models \text{Th}_{\forall^-}(\mathfrak{A})$ and $h : \mathfrak{A} \rightarrow \mathfrak{B}$ is a homomorphism. Suppose $\mathfrak{B} \models \varphi(h(\bar{a}))$, where $\varphi(\bar{x})$ is an ep-formula and \bar{a} is a tuple from A ; we must prove that $\mathfrak{A} \models \varphi(\bar{a})$. Let $\langle a_\alpha \rangle_{\alpha < |A|}$ well-order A . Consider $(\mathfrak{B}; \langle h(a_\alpha) \rangle_{\alpha < |A|})$, by Löwenheim-Skolem there is an elementarily equivalent $(\mathfrak{B}'; \langle h(a_\alpha) \rangle_{\alpha < |A|})$ such that \mathfrak{B}' is of cardinality no greater than $|A|$, $\mathfrak{B}' \models \text{Th}_{\forall^-}(\mathfrak{A})$, $h : \mathfrak{A} \rightarrow \mathfrak{B}'$ is a homomorphism and $\mathfrak{B}' \models \varphi(h(\bar{a}))$. Since $\mathfrak{B}' \models \text{Th}_{\forall^-}(\mathfrak{A})$ there is a homomorphism $g : \mathfrak{B}' \rightarrow \mathfrak{A}$ by Lemma 2.1 and ep-saturation of \mathfrak{A} . Therefore $\mathfrak{A} \models \varphi(g \circ h(\bar{a}))$ where $g \circ h$ is an endomorphism of \mathfrak{A} , which must be an embedding since \mathfrak{A} is a core. The result $\mathfrak{A} \models \varphi(\bar{a})$ follows as \mathfrak{A} is expanded by all ep-definable relations (note that ep-definable relations are not in general preserved by inverse embeddings, but atomic relations are). \square

We note that the additional assumption of being expanded by all ep-definable relations was necessary. Let \mathbb{Q}_0^+ be the set $\{q \in \mathbb{Q} : q \geq 0\}$. Then $(\mathbb{Q}_0^+; <)$ is a saturated core but is not epc for $\text{Th}_{\forall^-}(\mathbb{Q}_0^+; <)$.

4.2. Core theories. We now introduce the notion of core theories, models of which will be cores. We develop a certain existential-positive model theory, akin to the existential model theory developed in [32, 33]. Ultimately this will enable us to give a necessary and sufficient condition as to when a template \mathfrak{A} has an equivalent \mathfrak{B} (i.e., s.t. $\text{CSP}(\mathfrak{A}) = \text{CSP}(\mathfrak{B})$) such that \mathfrak{B} is a model-complete core.

We need the concept of a *diagram* of a structure, in various variants. For a model \mathfrak{A} , let \mathfrak{A}^c be the expansion of \mathfrak{A} by $|A|$ constants naming the elements of \mathfrak{A} . Let $\text{diag}_{\text{fo}}(\mathfrak{A})$ denote the *elementary diagram* of \mathfrak{A} , that is, the set of all fo-sentences true in the structure \mathfrak{A}^c . The set of all qf-sentences that hold on \mathfrak{A}^c is denoted by $\text{diag}(\mathfrak{A})$, and the set of all universal negative sentences that hold on \mathfrak{A}^c is denoted by $\text{diag}_{\forall^-}(\mathfrak{A})$. Finally, $\text{diag}_+(\mathfrak{A})$ denotes the set of all positive qf-sentences that hold on \mathfrak{A}^c .

Lemma 4.8 (Diagram lemma; Lemma 1.4.2. in [22]). *Let \mathfrak{A} and \mathfrak{B} be τ -structures, and let C be the set of new constants involved in the expansion \mathfrak{A}^c . Then the following are equivalent.*

- (1) *There is a $\tau \cup C$ -expansion \mathfrak{B}' of \mathfrak{B} s.t. $\mathfrak{B}' \models \text{diag}_+(\mathfrak{A})$.*
- (2) *There is a homomorphism from \mathfrak{A} to \mathfrak{B} .*

We will make use of the following well-known preservation theorems.

Theorem 4.9 (Homomorphism Preservation Theorem; see e.g. Exercise 2 in Section 5.5 of [22]). *Let T be a fo-theory. A fo-formula φ is equivalent to an ep-formula modulo T if and only if φ is preserved by all homomorphisms between models of T .*

Note that ep and pp are not interchangeable in the previous theorem.

Theorem 4.10 (Chang-Łoś-Suszko Theorem; Theorem 6.5.9 in [21] and remarks after the proof). *Let T be an fo- τ -theory. An fo- τ -formula φ is equivalent modulo T to a $\forall\exists$ -formula if and only if φ is preserved in unions of chains $(A_i)_{i \leq \kappa}$ whenever $\bigcup A_i$ and all the A_i are models of T .*

4.2.1. Positive version of Chang-Łoś-Suszko. The following fact is called the *existential amalgamation theorem* in [22].

Proposition 4.11 (Theorem 5.4.1 in [22]). *Let \mathfrak{A} and \mathfrak{B} be τ -structures, \bar{a} a sequence of elements of \mathfrak{A} , and $h : \bar{a} \rightarrow \mathfrak{B}$ an embedding such that every existential sentence true in $(\mathfrak{B}, h(\bar{a}))$ is also true in (\mathfrak{A}, \bar{a}) . Then there exists an elementary extension \mathfrak{C} of \mathfrak{A} and an embedding $g : \mathfrak{B} \rightarrow \mathfrak{C}$ such that $g(h(\bar{a})) = \bar{a}$.*

The existential amalgamation theorem has the following positive variant. We give the proof (it is omitted in [22]).

Proposition 4.12 (Theorem 5.4.7 in [22]). *Let \mathfrak{A} and \mathfrak{B} be τ -structures, \bar{a} a sequence of elements of \mathfrak{A} , and $h : \bar{a} \rightarrow \mathfrak{B}$ an embedding such that every existential positive sentence true in $(\mathfrak{B}, h(\bar{a}))$ is also true in (\mathfrak{A}, \bar{a}) . Then there exists an elementary extension \mathfrak{C} of \mathfrak{A} and a homomorphism $g : \mathfrak{B} \rightarrow \mathfrak{C}$ such that $g(h(\bar{a})) = \bar{a}$.*

Proof. Similar to the proof of 5.4.1 in [22]. Since h is an embedding, we can replace \mathfrak{B} by an isomorphic copy and assume that h is the identity on \bar{a} . Let \mathfrak{A}' be the expansion of \mathfrak{A} by those constant symbols that denote the elements of \bar{a} in \mathfrak{B}^c . By Lemma 4.8, it suffices to show that the theory $T := \text{diag}_{\text{fo}}(\mathfrak{A}') \cup \text{diag}_+(\mathfrak{B})$ is consistent. If T has no model, then by the compactness theorem there is a conjunction $\varphi(\bar{a}, \bar{c})$ of finitely many sentences in $\text{diag}_+(\mathfrak{B})$ such that $\mathfrak{A} \models \neg \exists \bar{y}. \varphi(\bar{a}, \bar{y})$. Since φ is quantifier-free and positive, the assumptions imply that $\mathfrak{B} \models \neg \exists \bar{y}. \varphi(\bar{a}, \bar{y})$. This contradicts that $\varphi(\bar{a}, \bar{c})$ is true in \mathfrak{B} . \square

Lemma 4.13. *Let T be a theory and \mathfrak{A} a model of the $\forall\exists^+$ -consequences of T . Then \mathfrak{A} can be extended to a model \mathfrak{B} of T such that every ep-formula that holds on a tuple \bar{a} in $A \subseteq \mathfrak{B}$ also holds on \bar{a} in \mathfrak{A} .*

Proof. It suffices to prove that $T \cup \text{diag}(\mathfrak{A}) \cup \text{diag}_{\forall^-}(\mathfrak{A})$ has a model \mathfrak{B} . Suppose for contradiction that it were inconsistent; then by compactness, there exists a finite subset U of $\text{diag}_{\forall^-}(\mathfrak{A}) \cup \text{diag}(\mathfrak{A})$ such that $T \cup U$ is inconsistent. Let φ be the conjunction over U where all constants are existentially quantified. Then $T \cup \{\varphi\}$ is inconsistent as well, and $\neg\varphi$ is equivalent to a $\forall\exists^+$ sentence, and a consequence of T . Hence, $\mathfrak{A} \models \neg\varphi$, a contradiction. \square

The following is a positive version of the Chang-Łoś-Suszko theorem (Theorem 4.10). We say that an fo-formula $\varphi(x_1, \dots, x_n)$ is *preserved in direct limits of models of T* if for all sequences $(\mathfrak{A}_i)_{i < \kappa}$ where all the \mathfrak{A}_i and $\mathfrak{A} := \lim_{i < \kappa} \mathfrak{A}_i$ are models of T , and every n -tuple \bar{a} of elements of A we have $\mathfrak{A} \models \varphi(\bar{a})$ whenever there is for every $i < \kappa$ an n -tuple \bar{a}^i where the j -th entry is a representative of the j -th entry of \bar{a} , and $\mathfrak{A}_i \models \varphi(\bar{a}^i)$.

Proposition 4.14. *Let T be a τ -theory, and Φ a set of τ -formulas (in a finite set of free variables). Then the following are equivalent.*

- (1) Φ is modulo T equivalent to a set of $\forall\exists^+$ -formulas;
- (2) Φ is preserved in direct limits of sequences of models of T ;
- (3) Φ is preserved in direct limits of countable sequences of models of T .

Proof. (1) implies (2) follows from Proposition 4.3.

The implication from (2) to (3) is trivial. For the implication from (3) to (1), assume that Φ is preserved by direct limits of sequences (\mathfrak{A}_i) as in the statement of the proposition. We can assume that Φ is a set of sentences (by adding constants; see the proof of Theorem 5.4.4. in [22]). Let Ψ be the set of all $\forall\exists^+$ -sentences that are consequences of $T \cup \Phi$. We show that $T \cup \Psi$ implies Φ , whereupon the result follows. It suffices to show that every model of $T \cup \Psi$ is elementary equivalent to a direct limit of a sequence $(\mathfrak{B}_i)_{i < \omega}$ of models of $T \cup \Phi$ where there are homomorphisms $f_{ij} : \mathfrak{B}_i \rightarrow \mathfrak{B}_j$ with $f_{jk} \circ f_{ij} = f_{ik}$ for all $i \leq j \leq k$.

To construct this sequence, we define an elementary chain of models $(\mathfrak{A}_i)_{i < \omega}$ of $T \cup \Psi$ such that there are

- homomorphisms $f_i : \mathfrak{A}_i \rightarrow \mathfrak{B}_i$, with $\mathfrak{B}_i \models T \cup \Phi$, such that for every tuple \bar{a}_i of elements from \mathfrak{A}_i and every existential positive formula θ , if $\mathfrak{B}_i \models \theta(f_i(\bar{a}_i))$, then $\mathfrak{A}_i \models \theta(\bar{a}_i)$, and
- homomorphisms $g_i : \mathfrak{B}_i \rightarrow \mathfrak{A}_{i+1}$, such that $g_i \circ f_i$ is the identity on \mathfrak{A}_i .

Let \mathfrak{A}_0 be a countable model of $T \cup \Psi$. To construct the rest of the sequence, suppose that \mathfrak{A}_i has been chosen.

Since \mathfrak{A}_0 is an elementary substructure of \mathfrak{A}_i , in particular all the $\forall\exists^+$ -consequences of $T \cup \Phi$ hold in \mathfrak{A}_i . By Lemma 4.13, the structure \mathfrak{A}_i can be extended to a model \mathfrak{B}_i of $T \cup \Phi$ such that every ep-sentence that holds in $(\mathfrak{B}_i, \bar{a}_i)$ also holds in $(\mathfrak{A}_i, \bar{a}_i)$ (i.e. f_i is the identity). Then by Proposition 4.12, there are an elementary extension \mathfrak{A}_{i+1} of \mathfrak{A}_i and a homomorphism $g_i : \mathfrak{B}_i \rightarrow \mathfrak{A}_{i+1}$ such that $g_i \circ f_i$ is the identity on \mathfrak{A}_i .

Then $\mathfrak{C} := \bigcup_{i < \omega} \mathfrak{A}_i$ equals $\lim_{i < \omega} \mathfrak{B}_i$, and by the Tarski-Vaught elementary chain theorem (Theorem 2.5.2 in [22]) \mathfrak{A}_0 is an elementary substructure of \mathfrak{C} . So \mathfrak{C} is a model of T , and the direct limit of models \mathfrak{B}_i of $T \cup \Phi$, and hence $\mathfrak{C} \models \Phi$. \square

Note that the variant of Proposition 4.14, in which Φ is a single formula φ , holds as its corollary, via an application of compactness at the end.

4.2.2. Model-complete core theories. A theory T is *model-complete* if any embedding across models of T is elementary, and T is called a *core theory* if every homomorphism between models of T is an embedding. The theory of a core need not be a core theory. The integers with successor $(\mathbb{Z}; \text{succ})$ is a core, but the homomorphisms from the non-atomic countable models of its theory into $(\mathbb{Z}; \text{succ})$ are not even injective (and so are not embeddings). The following closely follows the proof of Theorem 8.3.1 in [21].

Theorem 4.15. *Let T be an fo-theory over signature τ . Then the following are equivalent.*

- (1) *T is a model-complete core theory.*
- (2) *Every model of T is a model epc for T .*
- (3) *If $\mathfrak{A}, \mathfrak{B}$ are models of T and h is a homomorphism from \mathfrak{A} to \mathfrak{B} then there are an elementary extension \mathfrak{C} of \mathfrak{A} and an embedding g of \mathfrak{B} into \mathfrak{C} such that $g \circ h$ is the identity on \mathfrak{A} .*
- (4) *Every fo-formula is equivalent to an ep-formula modulo T .*

Proof. (1) implies (2) is immediate from the definition of epc models: if \mathfrak{A} and \mathfrak{B} are models of T and $h : A \rightarrow B$ is a homomorphism from \mathfrak{A} to \mathfrak{B} , then h must be an embedding since T is a core theory, and in fact must be elementary since T is model-complete. Hence, for every tuple \bar{a} from A and any ep-formula φ such that $h(\bar{a})$ satisfies φ we have that \bar{a} also satisfies φ .

(2) implies (3). Assume (2). Let \mathfrak{A} and \mathfrak{B} be models of T , and let h be a homomorphism from \mathfrak{A} to \mathfrak{B} . Choose \bar{a} to be a vector that enumerates the elements of \mathfrak{A} . Since \mathfrak{A} is epc for T , h is an embedding. Hence, every existential sentence that holds in $(\mathfrak{B}, h(\bar{a}'))$ also holds in (\mathfrak{A}, \bar{a}') , for any finite subtuple \bar{a}' of \bar{a} . Proposition 4.11 now directly implies (3).

(3) implies (4). We first claim that if (3) holds, then every homomorphism between models of T preserves all universal τ -formulas. For if h is a homomorphism of \mathfrak{A} into \mathfrak{B} , \bar{a} a tuple from A and $\varphi(\bar{x})$ a universal τ -formula such that $\mathfrak{A} \models \varphi(\bar{a})$, then taking \mathfrak{C} and g as in (3) we have $\mathfrak{C} \models \varphi(g(h(\bar{a})))$ and so $\mathfrak{B} \models \varphi(h(\bar{a}))$ since φ is a universal formula. This proves the claim. It follows by Theorem 4.9 that all universal τ -formulas are equivalent to ep- τ -formulas.

To finally prove (4), let $\varphi(\bar{x})$ be any fo- τ -formula, w.l.o.g. in prenex normal form. By a simple induction on the number of quantifier-blocks we can transform φ to an existential formula, using the fact that the innermost quantifier block is either existential or universal, and can therefore be transformed into an existential formula (see Theorem 8.3.1 in [21]). Finally, existential τ -formulas are preserved by homomorphisms between models of T , since such homomorphisms must be embeddings. Hence, the entire formula is even equivalent to an ep-formula by Theorem 4.9.

(4) implies (1). Let \mathfrak{A} and \mathfrak{B} be models of T . Any homomorphism from \mathfrak{A} to \mathfrak{B} preserves all ep-formulas and therefore all fo-formulas, and hence must be an elementary embedding. \square

From this we obtain the following. The proof is similar to the one for Lindström's test (see Theorem 8.3.4 in [21]).

Proposition 4.16 (Positive version of Lindström's test). *Let T be a theory with signature τ which has no finite models but for which there exists a model epc for T of cardinality $\lambda \geq |\tau|$. If T is λ -categorical, then T is a model-complete core theory.*

Sketch. We prove that every model of T is a model epc for T and use Theorem 4.15. So let \mathfrak{A} and \mathfrak{B} be two models of T and let h be a homomorphism from \mathfrak{A} to \mathfrak{B} . Let \bar{a} be a tuple such that $\mathfrak{B} \models \varphi(h(\bar{a}))$ and suppose for contradiction that $\mathfrak{A} \not\models \varphi(\bar{a})$. Then we can put those two structures into a new 2-sorted structure (comprising \mathfrak{A} , \mathfrak{B} and the homomorphism h between them) with fo-theory T^+ , and apply the Löwenheim-Skolem theorem to produce a countable model of T^+ (where both sorts have the same cardinality since T has no finite models). By applying Löwenheim-Skolem again, this time to T^+ augmented by sentences expressing a bijection between the two sorts (over a signature expanded by a new function symbol), we obtain a two-sorted model of T^+ where each sort has cardinality λ , inducing structures \mathfrak{C} and \mathfrak{D} , respectively. By assumption there exists a model epc for T of cardinality λ , with $\lambda \geq |\tau|$, and by λ -categoricity \mathfrak{C} is a model epc for T . This contradicts the fact that we can express in T^+ that \mathfrak{A} is not a model epc for T . \square

Proposition 4.17. *Let T be a model-complete core theory. Then T is equivalent to a $\forall\exists^+$ -theory.*

Proof. This is an immediate consequence of Proposition 4.14, since for any sequence $(\mathfrak{B}_i)_{i < \kappa}$ of models of T with homomorphisms $g_{ij} : \mathfrak{B}_i \rightarrow \mathfrak{B}_j$, the g_{ij} are elementary. By the Tarski-Vaught theorem on unions of elementary chains (Theorem 2.5.2 in [22]), we have that $\lim_{i < \kappa} \mathfrak{B}_i \models T$. \square

Proposition 4.18. *A λ -categorical τ -structure \mathfrak{A} , with $\lambda \geq |\tau|$, and epc for the $\forall\exists^+$ restriction of $\text{Th}(\mathfrak{A})$, is a model-complete core iff it has a theory equivalent to a $\forall\exists^+$ -theory.*

Proof. The forwards direction follows from Proposition 4.17. For the backward direction, we use Proposition 4.16 to derive that $\text{Th}(\mathfrak{A})$ is a model-complete core theory, whereupon \mathfrak{A} is a model-complete core. \square

Proposition 4.19. *Let T be an ω -categorical theory without finite models (over a finite or countable signature). Then T is a model-complete core theory if and only if it is equivalent to a $\forall\exists^+$ -theory.*

Proof. The forward direction follows from Proposition 4.17. For the backward direction, note that there exists an at most countable epc model for T (which is ω -categorical). We now apply Proposition 4.16 to deduce that T is a model-complete core theory. \square

4.2.3. Core companions.

Definition 4.20. Let T be a fo- τ -theory. Then a τ -theory U is called a *core companion* of T if

- U is a model-complete core theory;
- every model of U homomorphically maps to a model of T ;
- every model of T homomorphically maps to a model of U .

Proposition 4.21. *Let T and T' be τ -theories. The following are equivalent.*

- (1) Every model of T has a homomorphism to a model of T' , and every model of T' has a homomorphism to a model of T .
- (2) T and T' entail the same universal negative sentences.

Proof. To prove the implication from (1) to (2), assume (1), and let φ be a universal negative sentence entailed by T' . Suppose for contradiction that T has a model \mathfrak{C} such that $\mathfrak{C} \models \neg\varphi$. By (1), there is a homomorphism from \mathfrak{C} to a model \mathfrak{B} of T' . Since $\neg\varphi$ is equivalent to an existential positive sentence, is preserved by homomorphisms, and hence we have a contradiction to the assumption that T' entails φ .

For the implication from (2) to (1), assume (2), and let \mathfrak{B} be a model of T . Let S be the ep-theory of \mathfrak{B} . We claim that $S \cup T'$ is satisfiable. If not, then by compactness there is some finite subset $\{\varphi_1, \dots, \varphi_k\}$ of S such that T' entails $(\neg\varphi_1 \vee \dots \vee \neg\varphi_k)$. The formula $\neg\varphi_1 \vee \dots \vee \neg\varphi_k$ is equivalent to a universal negative sentence ψ , and T' entails ψ , so by (2) we have that T entails ψ , and hence $\mathfrak{B} \models \psi$. We have reached a contradiction, since $\mathfrak{B} \models \varphi_i$ for all $i \leq k$. So there indeed exists a model \mathfrak{A} of $S \cup T'$. The positive variant of the existential amalgamation theorem (Proposition 4.12) applied to \mathfrak{A} and \mathfrak{B} for the empty sequence \bar{a} gives a model \mathfrak{C} of $T' \cup S$ and a homomorphism from \mathfrak{B} to \mathfrak{C} . \square

The following two propositions are similar to Theorem 8.3.6 in [21].

Proposition 4.22. *Let T be a $\forall\exists^+$ -theory with signature τ . If T has a core companion U , then U is up to equivalence of theories unique, and is the theory of the class of all models epc for T .*

Proof. We show that the models epc for T are precisely the models of U . First assume that \mathfrak{A} is a model of U . Since U is a core companion of T , there is a homomorphism e from \mathfrak{A} to a model \mathfrak{B} of T . The assumption that U is a core companion of T also implies that there exists a homomorphism f from \mathfrak{B} into a model \mathfrak{C} of U . Then $f \circ e$ is a homomorphism between two models of U , and since U is a model-complete core theory it must be an elementary embedding. This shows in particular that e is an embedding. We claim that \mathfrak{A} is a model of the $\forall\exists^+$ -theory T . Let $\varphi = \forall\bar{y}.\psi$ be a sentence from T where ψ is a disjunction of ep and negated atomic τ -formulas, and let \bar{a} be a tuple from \mathfrak{A} . Since \mathfrak{B} is a model of T and therefore satisfies $\forall\bar{y}.\psi$, in particular the tuple $e(\bar{a})$ satisfies ψ . If $e(\bar{a})$ satisfies a negated atom in the disjunction ψ then \bar{a} satisfies ψ as e is an embedding. Otherwise, $e(\bar{a})$ satisfies an ep-formula in the disjunction ψ and $f \circ e(\bar{a})$ satisfies ψ follows, as f is a homomorphism. Now \bar{a} satisfies ψ as $f \circ e$ is an elementary embedding. Since this holds for all \bar{a} , we have proven that \mathfrak{A} satisfies φ .

In fact, \mathfrak{A} is a model epc for T . To verify this, let g be a homomorphism from \mathfrak{A} into another model \mathfrak{D} of T , \bar{a} a tuple from A , and φ an ep-formula with $\mathfrak{D} \models \varphi(g(\bar{a}))$. We have to show that $\mathfrak{A} \models \varphi(\bar{a})$. Again, since U is a core companion of T there exists a homomorphism h from \mathfrak{D} into a model \mathfrak{E} of U . Since U is a model-complete core theory, the mapping $h \circ g$ is elementary. Since h is a homomorphism, $\mathfrak{E} \models \varphi(h(g(\bar{a})))$. Since $h \circ g$ is elementary, $\mathfrak{A} \models \varphi(\bar{a})$.

Conversely, we show that every model \mathfrak{A} epc for T satisfies U . Since U is a core companion, there is a homomorphism h from \mathfrak{A} to a model \mathfrak{B} of U . By Proposition 4.17, U is equivalent to a $\forall\exists^+$ -theory, and thus it suffices to show that \mathfrak{A} satisfies all $\forall\exists^+$ -consequences $\forall\bar{y}.\psi(\bar{y})$ of U , where ψ is a disjunction of ep and negated atomic τ -formulas. Let \bar{a} be a tuple of elements of \mathfrak{A} . We have to show that $\mathfrak{A} \models \psi(\bar{a})$. Since $\mathfrak{B} \models \forall\bar{y}.\psi(\bar{y})$, at least one disjunct $\theta(h(\bar{a}))$ of $\psi(h(\bar{a}))$ is true in \mathfrak{B} . If θ is a negated atomic formula, then

$\theta(\bar{a})$ is also true in \mathfrak{A} since h is a homomorphism. If θ is an ep-formula, then we deduce \mathfrak{C} , a model of T , and a homomorphism $g : \mathfrak{B} \rightarrow \mathfrak{C}$ such that $\mathfrak{C} \models \theta(g(h(\bar{a})))$. Now $\mathfrak{A} \models \theta(\bar{a})$ since \mathfrak{A} is epc for T . In both cases we can conclude that $\mathfrak{A} \models \psi(\bar{a})$. \square

Proposition 4.23. *Let T be a $\forall\exists^+$ -theory with signature τ . Then T has a core companion if and only if the class of models epc for T is axiomatizable by a τ -theory.*

Proof. If T has a core companion U , then Proposition 4.22 above implies that U axiomatizes the class of models epc for T .

For the converse, suppose that the class of models epc for T is the class of all models of a τ -theory U . Then every model of U is in particular a model of T , and every model of T homomorphically maps to a model epc for T (and model of U) by Proposition 4.5. So we only have to verify that U is a model-complete core theory to show that U is the core companion of T .

Every model \mathfrak{A} of U is a model epc for T . Since being a model of U implies being a model of T , we deduce that \mathfrak{A} is in fact a model epc for U . It follows by the equivalence of (1) and (2) in Theorem 4.15 that U is a model-complete core theory. \square

Corollary 4.24. *Let \mathfrak{A} be a τ -structure, then there exists a τ -structure \mathfrak{B} such that $\text{CSP}(\mathfrak{A}) = \text{CSP}(\mathfrak{B})$ and \mathfrak{B} has a model-complete core theory iff the class of models epc for the $\forall\exists^+$ -theory of \mathfrak{A} is axiomatizable by a τ -theory.*

Proof. (Forwards.) If such a \mathfrak{B} exists then $\text{Th}(\mathfrak{B})$ must be a model-complete core theory and $\text{Th}(\mathfrak{B})$ is the core companion of $\text{Th}(\mathfrak{A})$. The result follows from Proposition 4.23.

(Backwards.) If the class of models epc for the $\forall\exists^+$ -theory of \mathfrak{A} is axiomatizable, then the $\forall\exists^+$ -theory of \mathfrak{A} has a core companion U . Let \mathfrak{B} be a model of U . It follows from Definition 4.20 that \mathfrak{B} is a model-complete core and that \mathfrak{A} and \mathfrak{B} are homomorphically equivalent, hence $\text{CSP}(\mathfrak{A}) = \text{CSP}(\mathfrak{B})$. \square

Note that \mathfrak{A} and \mathfrak{A}' may generate the same CSP, i.e. share the same ep-theory, whilst having distinct $\forall\exists^+$ -theories. The previous proof shows us that ep-equivalence does not affect the question as to whether or not the class of models epc for the $\forall\exists^+$ -theory of \mathfrak{A} is axiomatizable.

4.3. Equivalent ω -categorical templates. We are now in a position to establish precisely when a template \mathfrak{A} has an equivalent \mathfrak{B} (i.e., s.t. $\text{CSP}(\mathfrak{A}) = \text{CSP}(\mathfrak{B})$) such that \mathfrak{B} is ω -categorical. At the same time we reprove the result of [3] that every ω -categorical template is equivalent to a unique template that is a model-complete core. The result in [3] was obtained using ad-hoc arguments that poorly reflected the aspects of our programme for the existential-positive that are similar to the programme for the existential in [32, 33].

A structure is *homogeneous* (sometimes called *ultrahomogeneous* [21]) if every finite partial automorphism can be extended to a full automorphism.

Lemma 4.25. *Let \mathfrak{A} be a countable homogeneous structure such that for each k only a finite number of distinct k -ary relations can be defined by atomic formulas. Then \mathfrak{A} is ω -categorical.*

Proof. By homogeneity of \mathfrak{A} , the atomic formulas that hold on the elements of t in \mathfrak{A} determine the orbit of t under $\text{Aut}(\mathfrak{A})$. Since there are only finitely many such atomic formulas, it follows that there are finitely many orbits of k -tuples in $\text{Aut}(\mathfrak{A})$. The claim follows by Theorem 2.2. \square

A theory T is said to have the *Joint Homomorphism Property* (JHP) if, for all models \mathfrak{A} and \mathfrak{B} of T , there exists a model \mathfrak{C} of T such that both \mathfrak{A} and \mathfrak{B} map homomorphically into \mathfrak{C} .

Proposition 4.26. *For any theory T , the following are equivalent.*

- (1) *For all ep- τ -sentences φ_1 and φ_2 , if $T \cup \{\varphi_1\}$ is satisfiable and $T \cup \{\varphi_2\}$ is satisfiable then $T \cup \{\varphi_1, \varphi_2\}$ is satisfiable as well.*
- (2) *T has the JHP.*

Proof. Assume (1), and let τ be the signature of T . Let \mathfrak{A} and \mathfrak{B} be models of T . We claim that the theory $T' := T \cup \text{diag}_+(\mathfrak{A}) \cup \text{diag}_+(\mathfrak{B})$ is satisfiable. By compactness, it suffices to show that every finite subset S of T' is satisfiable. Let $S_1 := S \cap \text{diag}_+(\mathfrak{A})$ and $S_2 := S \cap \text{diag}_+(\mathfrak{B})$. By forming a finite conjunction, we see that S_1 and S_2 are logically equivalent to single sentences φ_1 and φ_2 , respectively. Let φ'_1 and φ'_2 be the τ -sentences corresponding to φ_1 and φ_2 , respectively, in which the new constant symbols are existentially quantified out. Certainly $T \cup \{\varphi'_1\}$ and $T \cup \{\varphi'_2\}$ are satisfiable since \mathfrak{A} and \mathfrak{B} are models of T and therefore satisfy all sentences from T . By (1), the theory $T \cup \{\varphi'_1, \varphi'_2\}$ is satisfiable as well. Therefore the claim is true, and there exists a model \mathfrak{C}' of T' . Let \mathfrak{C} be the τ -reduct of \mathfrak{C}' . Finally, Lemma 4.8 asserts the existence of a homomorphism from \mathfrak{A} to \mathfrak{C} , and from \mathfrak{B} to \mathfrak{C} , which proves (2).

For the implication from (2) to (1), suppose that T is a τ -theory with the JHP, and that φ_1 and φ_2 are ep-sentences such that $T \cup \{\varphi_1\}$ has a model \mathfrak{A} and $T \cup \{\varphi_2\}$ has a model \mathfrak{B} . By (2), there exists a model of T such that \mathfrak{A} and \mathfrak{B} homomorphically map to \mathfrak{C} . Then \mathfrak{C} clearly satisfies $T \cup \{\varphi_1, \varphi_2\}$. \square

For a satisfiable theory T , let \sim_n^T be the equivalence relation defined on ep-formulas with n free variables x_1, \dots, x_n as follows. For two such formulas φ_1 and φ_2 , let $\varphi_1 \sim_n^T \varphi_2$ iff for all ep-formulas ψ with free variables x_1, \dots, x_n we have that $\{\varphi_1, \psi\} \cup T$ is satisfiable if and only if $\{\varphi_2, \psi\} \cup T$ is satisfiable. By proving that a model epc for a certain type of theory is in fact ω -categorical, we will derive the following.

Theorem 4.27. *Let T be a theory with the JHP. Then the following are equivalent.*

- (i) *T has a core companion, unique up to equivalence of theories, that is either ω -categorical or the theory of a finite structure.*
- (ii) *\sim_n^T has finite index for each n .*
- (iii) *T has finitely many maximal ep- n -types for each n .*
- (iv) *There is a finite or ω -categorical model \mathfrak{A} that is a model-complete core, and which satisfies an existential positive sentence φ iff $T \cup \{\varphi\}$ is satisfiable.*

Proof. (i) \Rightarrow (ii). Let U be the core companion of T , since they entail the same universal negative sentences, we can deduce – for ep-formulas ψ – that $U \cup \{\psi\}$ is satisfiable if and only if $T \cup \{\psi\}$ is satisfiable. It follows that the indices of \sim_n^U and \sim_n^T coincide.

Let A be a finite, or the countable, model of U . For a proof by contraposition, assume \sim_n^U has infinite index for some n . Let φ_1 and φ_2 be two ep-formulas from different equivalence classes of \sim_n^U . Hence, there is an ep-formula φ_3 with free variables x_1, \dots, x_n such that exactly one of the two formulas $\varphi_1 \wedge \varphi_3$ and $\varphi_2 \wedge \varphi_3$ is satisfiable relative to U . This shows that φ_1 and φ_2 define over \mathfrak{A} distinct relations, and therefore that \mathfrak{A} can not be ω -categorical (as Theorem 2.2 asserts that it has only a finite number of inequivalent first-order definable relations of arity n , and in particular only a finite number of inequivalent ep-definable relations of arity n).

(ii) \Rightarrow (iii). We show that every maximal ep- n -type p is determined completely by the \sim_n^T equivalence classes of the ep-formulas contained in p . Since there are finitely many such classes, the result follows. Let p and q be maximal ep- n -types s.t. for every $\varphi_1 \in p$, exists $\varphi'_1 \in q$ s.t. $\varphi_1 \sim_n^T \varphi'_1$ and for every $\varphi_2 \in q$, exists $\varphi'_2 \in p$ s.t. $\varphi_2 \sim_n^T \varphi'_2$. We aim to prove that $p = q$. If not then there exists, w.l.o.g., $\psi \in p$ s.t. $\psi \notin q$. Clearly, $T \cup p \cup \{\psi\}$ is satisfiable, and, since q is maximal, $T \cup q \cup \{\psi\}$ is not satisfiable. By compactness $T \cup \{\theta_q, \psi\}$ is not satisfiable for some finite conjunction θ_q of formulas from q . Now, $\theta_q \in q$ by maximality and there exists by assumption $\theta'_q \in p$ s.t. $\theta_q \sim_n^T \theta'_q$. By definition of \sim_n^T we deduce $T \cup \{\theta'_q, \psi\}$ satisfiable iff $T \cup \{\theta_q, \psi\}$ satisfiable. Since the latter is not satisfiable, we deduce that neither is the former, which yields the contradiction that $T \cup p \cup \{\psi\}$ is not satisfiable.

(iii) \Rightarrow (iv). We aim to show that

$$S := \{\varphi : \varphi \text{ ep-formula consistent with } T\} \cup \{\neg\varphi : \varphi \text{ ep-formula not consistent with } T\}$$

has a finite or ω -categorical model that is a model-complete core. Since T has the JHP, it follows by induction that S is finitely satisfiable, and thence by compactness that S is satisfiable. We claim that S has finitely many maximal ep- n -types for each n . It suffices to prove that if p is a maximal ep- n -type of S then it is also a maximal ep- n -type of T . If φ is an ep-formula, then $T \cup p \cup \{\varphi\}$ being satisfiable implies $S \cup p \cup \{\varphi\}$ is satisfiable (by induction through JHP and compactness), and the claim follows. Let the number of maximal ep- n -types of S , μ_n , be finite for all n . We will show that S has an ω -categorical model. We consider the signature τ' , which is the expansion of τ by μ_n relations of each arity n , corresponding to the maximal ep- n -types of S . Any model of S has a canonical (unique) expansion to a τ' -structure (by the new relation symbols labeling tuples that attain their type). Consider the canonical τ' -expansion \mathfrak{A}' of a countable or finite τ -model \mathfrak{A} epc for S , guaranteed to exist by Proposition 4.5. We will shortly prove that \mathfrak{A}' is homogeneous. From this it will follow that \mathfrak{A} is ω -categorical by Lemma 4.25 (there is only a finite number of inequivalent atomic formulas of each arity n , since the collapsings of formulas with higher arity are themselves already n -ary atoms of τ'), whereupon ω -categoricity is inherited by its τ -reduct \mathfrak{A} . Since \mathfrak{A} is epc for S , it can be seen to also be epc for $\text{Th}_{\forall-}(\mathfrak{A})$, and it follows from Proposition 4.7 that \mathfrak{A} is a core. Homogeneity enforces that each fo-formula is equivalent modulo $\text{Th}(\mathfrak{A}')$ to a Boolean combination of ep-formulas. However, because we expand by a finite number of maximal ep- n -types, each complement of an ep- n -type is again an ep- n -type, and so the two coincide. It follows that each fo-formula is equivalent modulo $\text{Th}(\mathfrak{A}')$ to an ep-formula, and model-completeness of $\text{Th}(\mathfrak{A}')$ follows from Theorem 4.15. It follows that $\mathfrak{A} \models \text{Th}(\mathfrak{A}')$ is a model-complete core.

It remains to prove that \mathfrak{A}' is homogeneous. An ep-formula $\varphi(\bar{x})$ is said to isolate a maximal ep- n -type $p(\bar{x})$ of S , if p is the only maximal ep- n -type of S of which φ is a member. If there is only a finite number of maximal ep- n -types of S , then it follows that each has an isolating formula (i.e. a formula that is in that maximal ep- n -type but in no other maximal ep- n -type). Let $f : (a_1, \dots, a_m) \mapsto (b_1, \dots, b_m)$ be a partial automorphism of \mathfrak{A}' (in the signature τ'). Let a' be an arbitrary element of A' . Consider the ep- n -types $p(x_1, \dots, x_m)$ of (a_1, \dots, a_m) and $q(x_1, \dots, x_m, y)$ of (a_1, \dots, a_m, a') in \mathfrak{A} . By Proposition 4.6, each of these types is maximal, and is isolated by the ep-formulas $\theta_p(x_1, \dots, x_m)$ and $\theta_q(x_1, \dots, x_m, y)$, respectively. Furthermore, the type of (b_1, \dots, b_m) in \mathfrak{A} is p (as the partial automorphism of \mathfrak{A}' respects the signature τ'). But now, since $\exists y. \theta_q(x_1, \dots, x_m, y) \in p$ (by maximality), we may deduce a b' s.t. $\mathfrak{A}' \models \theta_q(b_1, \dots, b_m, b')$ and consequently $\mathfrak{A}' \models q(b_1, \dots, b_m, b')$. It follows that $f' : (a_1, \dots, a_m, a') \mapsto (b_1, \dots, b_m, b')$ is a partial automorphism of \mathfrak{A}' (in

the signature τ'). A simple back-and-forth argument shows that we may extend to an automorphism of \mathfrak{A}' , and the result follows.

(iv) \Rightarrow (i). Let \mathfrak{A} be the finite or ω -categorical structure that is a model-complete core. We claim $\text{Th}(\mathfrak{A})$ is a core companion of T . This follows from the fact that T and $\text{Th}(\mathfrak{A})$ entail the same universal negative sentences, via Proposition 4.21. Observe that $\text{Th}(\mathfrak{A})$ is equivalent to a $\forall\exists^+$ theory, by Proposition 4.17. Being a model-complete core theory follows immediately if \mathfrak{A} is finite, and via Proposition 4.16 if \mathfrak{A} is ω -categorical. Uniqueness follows from Proposition 4.22, since the core companion of $\text{Th}(\mathfrak{A})$ is, up to equivalence of theories, unique (and this implies that $\text{Th}(\mathfrak{A})$ is already the unique core companion of T , up to equivalence of theories). \square

We end this section by drawing comparison between Theorem 4.27 and a variant that is existential (but not existential positive). A theory T is said to have the *Joint Embedding Property* (JEP) if, for all models \mathfrak{A} and \mathfrak{B} of T , there exists a model \mathfrak{C} of T such that both \mathfrak{A} and \mathfrak{B} embed into \mathfrak{C} . A *companion* to a theory T is a theory T' s.t. every model of T' can be embedded into a model of T and vice-versa. A *model companion* is a companion that is model-complete. A \exists_1 -type is defined exactly as an ep-type but with existential (as opposed to existential positive) formulas.

Theorem 4.28 ([33]). *Let T be a theory with the JEP. Then the following are equivalent.*²

- T has a model companion that is either ω -categorical or the theory of a finite structure.
- T has finitely many maximal \exists_1 - n -types for each n .

We also draw the reader's attention to the similarity between our Proposition 4.6 and Theorem 1.2(b) in [33].

5. APPLICATIONS

We give a series of applications of our results in the study of the structure and complexity of CSPs.

5.1. Essentially unary polymorphisms. We will begin by demonstrating that the power of infinitary polymorphisms can be greatly limited. The forthcoming three lemmas are well-known for finite domains (also for ω -categorical structures). They require a little care in the infinite case.

A function $f : A^\alpha \rightarrow A$ is *essentially unary* if there exist a $\beta < \alpha$ and $g : A \rightarrow A$ such that, for all $\bar{x} \in A^\alpha$, $f(\bar{x}) = g(x_\beta)$. For $\bar{x}, \bar{w} \in A^\alpha$, and $X \subseteq \alpha$, let $\bar{x}[\bar{x}_X/\bar{w}_X]$ be the tuple \bar{x} with each entry x_β , where $\beta \in X$, substituted by w_β .

Lemma 5.1. *A function $f : A^\alpha \rightarrow A$ is not essentially unary iff there exist two non-empty and disjoint $X, Y \subseteq \alpha$, such that both*

- exist $\bar{x}, \bar{w}, \bar{w}' \in A^\alpha$ s.t. $f(\bar{x}[\bar{x}_X/\bar{w}_X]) \neq f(\bar{x}[\bar{x}_X/\bar{w}'_X])$, and
- exist $\bar{y}, \bar{z}, \bar{z}' \in A^\alpha$ s.t. $f(\bar{y}[\bar{y}_Y/\bar{z}_Y]) \neq f(\bar{y}[\bar{y}_Y/\bar{z}'_Y])$.

²Note that Simmons uses a stronger notion of model companion than we have given here, but the statement can be seen to hold nonetheless.

Proof. We will benefit from the following local definition. A set $Z \subseteq \alpha$ is termed *good* if the following holds: for all $\bar{x}, \bar{w}, \bar{w}' \in A^\alpha$ we have $f(\bar{x}[\bar{x}_Z/\bar{w}_Z]) = f(\bar{x}[\bar{x}_Z/\bar{w}'_Z])$. If Z is not good, then we term it *bad*. Note that good sets are closed under union; i.e., if X and Y are both good, then so is $X \cup Y$. The contrapositive of the lemma is the assertion that f is essentially unary iff, for any two non-empty and disjoint $X, Y \subseteq \alpha$, at least one of X and Y is good.

(Backwards.) By contraposition. If f is essentially unary, then let β and g be s.t. $f(\bar{x}) = g(x_\beta)$. Now, take any two non-empty and disjoint $X, Y \subseteq \alpha$. At least one does not contain β , and it must be a good set.

(Forwards.) By contraposition. Assume that, for any two non-empty and disjoint $X, Y \subseteq \alpha$, at least one of X and Y is good. If there are no bad subsets of α , i.e. f is constant, then clearly f is essentially unary. Assume the existence of some bad set. We will derive the existence of a bad set of cardinality 1; for otherwise let Z be a minimal bad set (under the total lexicographical order on the 0 – 1 characteristic sequence of length α) of cardinality greater than 1. Let Z_1 and Z_2 be a non-trivial partition of Z . At least one of Z_1 and Z_2 must be good, by assumption. Hence the other must be bad (as good sets are closed under union, and $Z := Z_1 \cup Z_2$ is bad), contradicting minimality of Z . Let $Z = \{\beta\}$ be a minimal bad set. Set

$$g(x_\beta) := f(x_\beta^\alpha) = f(x_\beta, x_\beta, \dots),$$

i.e. each variable x_γ , $\gamma \leq \alpha$, is substituted by x_β (of course the choice of x_β as the variable here is not important). That $f(\bar{x}) = g(x_\beta)$ now follows from $\alpha \setminus \{\beta\}$ being a good set. \square

Lemma 5.2. *For all \mathfrak{A} , $(x = y \vee u = v) \in \langle \mathfrak{A} \rangle_{\text{pp}}$ iff $\langle \mathfrak{A} \rangle_{\text{pp}} = \langle \mathfrak{A} \rangle_{\text{ep}}$.*

Proof. The backward direction is trivial. We prove the forward direction. Our proof will be by simulation of the binary \vee . Take $\varphi \in \langle \mathfrak{A} \rangle_{\text{ep}}$ in prenex form; we will recursively remove disjunctions of the form

$$\psi_1(x_1, \dots, x_n, y_1, \dots, y_p) \vee \psi_2(x_1, \dots, x_n, z_1, \dots, z_q).$$

We may assume that each of ψ_1 and ψ_2 is alone satisfiable, for otherwise their disjunction is logically equivalent to just one of them. We will introduce new variables $x'_1, \dots, x'_n, y'_1, \dots, y'_p$ and $x''_1, \dots, x''_n, z'_1, \dots, z'_q$. It follows from [9] that there is a $\theta \in \langle (A; x = y \vee u = v) \rangle_{\text{pp}}$ such that $\theta \equiv$

$$\begin{aligned} (x'_1 = x_1 \wedge \dots \wedge x'_k = x_k \wedge y'_1 = y_1 \wedge \dots \wedge y'_p = y_p) \quad \vee \\ (x''_1 = x_1 \wedge \dots \wedge x''_k = x_k \wedge z'_1 = z_1 \wedge \dots \wedge z'_q = z_q). \end{aligned}$$

The disjunct $\psi_1 \vee \psi_2$ should be replaced with the following, in which the existential quantifiers should be read as all coming before the conjunction.

$$\begin{aligned} \exists x'_1, \dots, x'_n, y'_1, \dots, y'_p. \quad \psi_1(x'_1, \dots, x'_n, y'_1, \dots, y'_p) \wedge \\ \exists x''_1, \dots, x''_n, z'_1, \dots, z'_q. \quad \psi_2(x''_1, \dots, x''_n, z'_1, \dots, z'_q) \wedge \\ \theta(x_1, \dots, x_n, y_1, \dots, y_p, \\ x'_1, \dots, x'_n, y'_1, \dots, y'_p, \\ x''_1, \dots, x''_n, z'_1, \dots, z'_q) \end{aligned}$$

\square

Lemma 5.3. *Let \mathfrak{A} be such that $(u = v \vee x = y) \in \langle \mathfrak{A} \rangle_{\text{pp}}$. Then all (finitary and infinitary) polymorphisms of \mathfrak{A} are essentially unary.*

Proof. Let $P_4 := (u = v \vee x = y) \in \langle \mathfrak{A} \rangle_{\text{pp}}$. It follows from Lemma 3.4 that P_4 must be preserved by the polymorphisms of \mathfrak{A} . Suppose for contradiction that \mathfrak{A} has a polymorphism $f : M^\alpha \rightarrow M$ that is not essentially unary. From Lemma 5.1, we deduce non-empty and disjoint $X, Y \subseteq \alpha$, s.t. there exist $\bar{x}, \bar{w}, \bar{w}' \in A^\alpha$ with $f(\bar{x}[\bar{x}_X/\bar{w}_X]) \neq f(\bar{x}[\bar{x}_X/\bar{w}'_X])$ and $\bar{y}, \bar{z}, \bar{z}' \in A^\alpha$ with $f(\bar{y}[\bar{y}_Y/\bar{z}_Y]) \neq f(\bar{y}[\bar{y}_Y/\bar{z}'_Y])$. But, for each $\beta \in \alpha$,

$$P_4(\bar{x}[\bar{x}_X/\bar{w}_X]_\beta, \bar{x}[\bar{x}_X/\bar{w}'_X]_\beta, \bar{y}[\bar{y}_Y/\bar{z}_Y]_\beta, \bar{y}[\bar{y}_Y/\bar{z}'_Y]_\beta)$$

holds, by disjointness of X and Y , while

$$P_4(f(\bar{x}[\bar{x}_X/\bar{w}_X]), f(\bar{x}[\bar{x}_X/\bar{w}'_X]), f(\bar{y}[\bar{y}_Y/\bar{z}_Y]), f(\bar{y}[\bar{y}_Y/\bar{z}'_Y]))$$

does not. □

Proposition 5.4. *For all structures \mathfrak{A} , $\langle \mathfrak{A} \rangle_{\text{pp}} = \langle \mathfrak{A} \rangle_{\text{ep}}$ iff all ω -polymorphisms of all elementary extensions of \mathfrak{A} are essentially unary.*

Proof. (Forwards.) If $\langle \mathfrak{A} \rangle_{\text{pp}} = \langle \mathfrak{A} \rangle_{\text{ep}}$ then $(u = v \vee x = y) \in \langle \mathfrak{A} \rangle_{\text{pp}}$, and so $(u = v \vee x = y) \in \langle \mathfrak{A}' \rangle_{\text{pp}}$ for all $\mathfrak{A}' \succeq \mathfrak{A}$. The result follows from Lemma 5.3.

(Backwards.) If all ω -polymorphisms of all elementary extensions of \mathfrak{A} are essentially unary, then in particular this is true of the monster elementary extension \mathfrak{M} built as in Lemma 3.3. It follows from Theorem 3.5 that $(u = v \vee x = y) \in \langle \mathfrak{M} \rangle_{\text{pp}}$, which gives $\langle \mathfrak{M} \rangle_{\text{pp}} = \langle \mathfrak{M} \rangle_{\text{ep}}$ by Lemma 5.2. The result $\langle \mathfrak{A} \rangle_{\text{pp}} = \langle \mathfrak{A} \rangle_{\text{ep}}$ follows since $\mathfrak{A} \preceq \mathfrak{M}$. □

We are able to prove that the stipulation of elementary extension in Proposition 5.4 is necessary, by exhibiting a structure whose ω -polymorphisms include only projections but for which $(x = y \vee u = v)$ is not pp-definable. In the following, $+$ should be read as a ternary relation.

Lemma 5.5. *The only ω -polymorphisms of $(\mathbb{Q}; +, 1, \neq)$ are projections.*

Proof. We give the proof for polymorphisms of arity ω , but the argument works just as well for any infinite or finite arity. A function $f : D^\omega \rightarrow D$ is *idempotent* if $f(d, d, \dots) = d$, for all $d \in D$. It is *conservative* if it further satisfies $f(d_1, d_2, \dots) \in \{d_1, d_2, \dots\}$, for all $d_1, d_2, \dots \in D$. Let $f : \mathbb{Q}^\omega \rightarrow \mathbb{Q}$ be a polymorphism of $(\mathbb{Q}; +, 1, \neq)$. It is clear that f is idempotent as the only endomorphism of $(\mathbb{Q}; +, 1)$ is the identity. Further, by preservation of \neq , it is easy to see that f must be conservative. Consider $\{0, 1\}^\omega$ with the total lexicographical ordering induced by $0 < 1$. Choose some minimal $\langle z_\lambda \rangle_{\lambda < \omega} \in \{0, 1\}^\omega$ s.t. $f(\langle z_\lambda \rangle_{\lambda < \omega}) = 1$ (since $f(1, 1, \dots) = 1$, such a $\langle z_\lambda \rangle_{\lambda < \omega}$ exists). If $\langle z_\lambda \rangle_{\lambda < \omega}$ had more than one index that is a 1, then there would exist $\langle z'_\lambda \rangle_{\lambda < \omega}$ and $\langle z''_\lambda \rangle_{\lambda < \omega}$ s.t. $\langle z'_\lambda \rangle_{\lambda < \omega}, \langle z''_\lambda \rangle_{\lambda < \omega} < \langle z_\lambda \rangle_{\lambda < \omega}$ and $\langle z'_\lambda \rangle_{\lambda < \omega} + \langle z''_\lambda \rangle_{\lambda < \omega} = \langle z_\lambda \rangle_{\lambda < \omega}$ and so, by preservation of $+$, one of $\langle z'_\lambda \rangle_{\lambda < \omega}, \langle z''_\lambda \rangle_{\lambda < \omega} = 1$, contradicting minimality of $\langle z_\lambda \rangle_{\lambda < \omega}$. So, for some i , $\langle z_\lambda \rangle_{\lambda < \omega}$ is of the form

$$(0, \dots, 0, \overbrace{1}^{\textit{ith position}}, 0, \dots).$$

By preservation of $+$, it follows, for each $q \in \mathbb{Q}$, that $f(q \cdot \langle z_\lambda \rangle_{\lambda < \omega}) = q$.

Firstly, we consider $\langle x_\lambda \rangle_{\lambda < \omega} \in \mathbb{Q}^\omega$ s.t. $q \notin \{x_\lambda : \lambda < \omega\} \neq \mathbb{Q}$. If $f(\langle x_\lambda \rangle_{\lambda < \omega}) = p \neq x_i$, then, by preservation of $+$,

$$f(\langle x_\lambda \rangle_{\lambda < \omega} + (q - p)\langle z_\lambda \rangle_{\lambda < \omega}) = f(\langle x_\lambda \rangle_{\lambda < \omega}) + (q - p)f(\langle z_\lambda \rangle_{\lambda < \omega}) = q.$$

But this violates conservativity of f as q does not appear in $\langle x_\lambda \rangle_{\lambda < \omega} + (q-p)\langle z_\lambda \rangle_{\lambda < \omega}$ (since $p \neq x_i$).

Finally, we take an arbitrary $\langle x_\lambda \rangle_{\lambda < \omega} \in \mathbb{Q}^\omega$. Consider the set $\Lambda := \{\lambda : x_\lambda = 1\}$ and $\langle x'_\lambda \rangle_{\lambda < \omega}$ and $\langle x''_\lambda \rangle_{\lambda < \omega}$ obtained according to $x'_\lambda = x_\lambda$, if $\lambda \notin \Lambda$, and $= 0$ otherwise; and $x''_\lambda = 1$, if $\lambda \in \Lambda$, and $= 0$ otherwise. Clearly $\langle x_\lambda \rangle_{\lambda < \omega} = \langle x'_\lambda \rangle_{\lambda < \omega} + \langle x''_\lambda \rangle_{\lambda < \omega}$, and $\langle x'_\lambda \rangle_{\lambda < \omega}$ and $\langle x''_\lambda \rangle_{\lambda < \omega}$ satisfy the condition of the previous paragraph, i.e. that neither $\{x'_\lambda : \lambda < \omega\}$ nor $\{x''_\lambda : \lambda < \omega\}$ is \mathbb{Q} . The result follows by preservation of $+$. \square

It follows that $(x = y \vee u = v) \in \text{Inv}(\text{Pol}^\omega(\mathbb{Q}; +, 1, \neq))$, though $(x = y \vee u = v)$ is not pp-definable in $(\mathbb{Q}; +, 1, \neq)$ since if it were we could also derive $(x = y \vee u = v) \in \langle (\mathbb{R}; +, 1, \neq) \rangle_{\text{pp}}$ (since $(\mathbb{R}; +, 1, \neq)$ and $(\mathbb{Q}; +, 1, \neq)$ share the same theory). This would contradict Lemma 5.3 as $(\mathbb{R}; +, 1, \neq)$ has polymorphisms that are not essentially unary: indeed, there is an isomorphism between $(\mathbb{R}; +, 1)^2$ and $(\mathbb{R}; +, 1)$ (that we shall use again shortly), which gives a bijective homomorphism from $(\mathbb{R}; +, 1, \neq)^2$ to $(\mathbb{R}; +, 1, \neq)$.

The following definition comes from [6]. For a structure \mathfrak{A} and an ep-sentence φ , we generate the boolean sentence $F_{\mathfrak{A}}(\varphi)$ by removing all existential quantifiers and replacing each atom $R(x_1, \dots, x_k)$, where $R^{\mathfrak{A}}$ is empty, with *false*, and replacing all other atoms with *true*. \mathfrak{A} is said to be *locally refutable* if for every ep-sentence φ , $\mathfrak{A} \models \varphi$ iff $F_{\mathfrak{A}}(\varphi)$ is true.

Proposition 5.6. *Let \mathfrak{A} be a structure that is not locally refutable and for which all ω -polymorphisms in all elementary extensions are essentially unary. Then $\text{CSP}(\mathfrak{A})$ is NP-hard.*

Proof. It is proved in [6] that the evaluation of ep-sentences on \mathfrak{A} is NP-hard. The result now follows from Lemma 5.2 (note that the recursive removal of disjunction induces a polynomial time reduction). \square

5.2. First-order definable CSPs. Recall $\varphi[\mathfrak{B}]$ to be the canonical query of \mathfrak{B} . $\text{CSP}(\mathfrak{A})$ is said to be *first-order definable* if there is an fo-sentence $\psi_{\mathfrak{A}}$ such that, for all finite \mathfrak{B} , $\mathfrak{A} \models \varphi[\mathfrak{B}]$ (i.e. $\varphi[\mathfrak{B}] \in \text{CSP}(\mathfrak{A})$) iff $\mathfrak{B} \models \psi_{\mathfrak{A}}$. The following definition comes from [26]. The *one-tolerant n -th power* ${}^1\mathfrak{A}^n$ of a τ -structure \mathfrak{A} is the τ -structure with domain A^n where a k -ary $R \in \tau$ denotes the relation consisting of all those k -tuples $((a_1^1, \dots, a_1^n), \dots, (a_k^1, \dots, a_k^n))$ such that

$$|\{j : (a_1^j, \dots, a_k^j) \in R^{\mathfrak{A}}\}| \geq n - 1.$$

For $n \geq 3$, an n -ary polymorphism f of \mathfrak{A} is called a *1-tolerant* polymorphism if f is a homomorphism from ${}^1\mathfrak{A}^n$ to \mathfrak{A} . The following is our analog of the result from [26].

Theorem 5.7. *Let \mathfrak{A} be a monster elementary extension (as constructed as in Lemma 3.3) on a finite signature. Then $\text{CSP}(\mathfrak{A})$ is first-order definable if and only if \mathfrak{A} has a 1-tolerant polymorphism.*

Proof. Claim 1. Let \mathfrak{M} be a monster extension. If all finite substructures \mathfrak{C} of ${}^1\mathfrak{M}^{n+1}$ map homomorphically to \mathfrak{M} , then ${}^1\mathfrak{M}^{n+1}$ maps homomorphically to \mathfrak{M} .

Proof of Claim 1. We note, for structures \mathfrak{A} and \mathfrak{B} , that if \mathfrak{A} and \mathfrak{B} are elementarily equivalent, then so are ${}^1\mathfrak{A}^{n+1}$ and ${}^1\mathfrak{B}^{n+1}$ (as ${}^1\mathfrak{A}^{n+1}$ is fo-definable in \mathfrak{A}). The assumption of the claim may be restated as that all pp- τ -sentences true in ${}^1\mathfrak{M}^{n+1}$ are true in \mathfrak{M} . This is clearly true by assumption that all finite substructures \mathfrak{C} map homomorphically to \mathfrak{M} .

A finite τ -structure \mathfrak{C} is an *obstruction* for the τ -structure \mathfrak{A} if there is no homomorphism from \mathfrak{C} to \mathfrak{A} . A family \mathcal{F} of obstructions for \mathfrak{A} is called a *complete set of obstructions*

if for every τ -structure \mathfrak{B} that does not admit a homomorphism to \mathfrak{A} there exists some $\mathfrak{C} \in \mathcal{F}$ which admits a homomorphism to \mathfrak{B} . The structure \mathfrak{A} is said to have *finite duality* if it admits a finite complete set of obstructions. An obstruction \mathfrak{C} for \mathfrak{A} is called *critical* if every proper (not necessarily induced) substructure of \mathfrak{C} admits a homomorphism to \mathfrak{A} . For any set A , let pr_k^n denote the projection map from A^n to A which maps any tuple to its k -th coordinate. We claim the following (essentially from [26]):

Claim 2. If there exists an $(n + 1)$ -ary 1-tolerant polymorphism of \mathfrak{A} then the critical obstructions of \mathfrak{A} have at most n hyperedges. If \mathfrak{A} is a monster extension, the converse holds as well.

Proof of Claim 2. (Forwards.) By contraposition. Let \mathfrak{C} be a critical obstruction of \mathfrak{A} with m distinct hyperedges e_1, \dots, e_m , $m > n$. Then for $k \in \{1, \dots, m\}$, the τ -structure \mathfrak{C}_k obtained from \mathfrak{C} by removing e_k (without changing the domain) admits a homomorphism h_k to \mathfrak{A} . By definition of ${}^1\mathfrak{A}^m$, the map $h = (h_1, \dots, h_m)$ is a homomorphism from \mathfrak{C} to ${}^1\mathfrak{A}^m$. Therefore there is no homomorphism from ${}^1\mathfrak{A}^m$ to \mathfrak{A} , and in particular none from ${}^1\mathfrak{A}^{n+1}$ to \mathfrak{A} .

(Backwards.) Conversely, suppose that \mathfrak{A} is a monster extension, and that there is no homomorphism from ${}^1\mathfrak{A}^{n+1}$ to \mathfrak{A} . It follows from Claim 1 that there exists a finite substructure \mathfrak{C} of ${}^1\mathfrak{A}^{n+1}$ which has no homomorphism to \mathfrak{A} . Hence, \mathfrak{C} is an obstruction of \mathfrak{A} which admits a homomorphism h to ${}^1\mathfrak{A}^{n+1}$. Let \mathfrak{C}' be a (not necessarily induced) substructure of \mathfrak{C} that is critical (such a \mathfrak{C}' always exists). For every $k \in \{1, \dots, n + 1\}$ there exists a hyperedge e_k of \mathfrak{C}' which is not preserved by $\text{pr}_k^{n+1} \circ h$, since $\text{pr}_k^{n+1} \circ h$ is not a homomorphism from \mathfrak{C} to \mathfrak{A} . By definition of ${}^1\mathfrak{A}^{n+1}$, e_k is respected by $\text{pr}_j^{n+1} \circ h$ for every $j \neq k$, and thus $e_j \neq e_k$ for $j \neq k$. Therefore \mathfrak{C} has at least $n + 1$ hyperedges. \square

Proof of Theorem. (Forwards.) Suppose first that $\text{CSP}(\mathfrak{A})$ is first-order definable. Since $\text{CSP}(\mathfrak{A})$ is a class of finite structures that is closed under inverse homomorphisms, by the dual version of Rossman's Theorem, [31], there is a universal negative first-order τ -sentence φ that holds on a finite structure \mathfrak{B} if and only if \mathfrak{B} homomorphically maps to \mathfrak{A} . Bringing φ into prenex negation normal form, it is straightforward to read from φ a finite complete set \mathcal{F} of obstructions to \mathfrak{A} . Let m be the maximal number of hyperedges in the obstructions from \mathcal{F} . By the claim above, since \mathfrak{A} is a monster, there is a homomorphism from ${}^1\mathfrak{A}^{m+1}$ to \mathfrak{A} . This is by definition a 1-tolerant polymorphism of \mathfrak{A} .

(Backwards.) Now suppose that for some n the structure ${}^1\mathfrak{A}^{n+1}$ admits a homomorphism to \mathfrak{A} . By the claim above the critical obstructions of \mathfrak{A} have at most n hyperedges. Since our signature is finite and relational, this implies that there are finitely many critical obstructions to \mathfrak{A} . This implies that the set of all critical obstructions is a finite obstruction set for \mathfrak{A} . It is now straightforward to write down a (universal) first-order definition of $\text{CSP}(\mathfrak{A})$. \square

Corollary 5.8. *Let \mathfrak{A} be a structure on a finite signature. Then $\text{CSP}(\mathfrak{A})$ is first-order definable if and only if \mathfrak{A} has an elementary extension which has a 1-tolerant polymorphism.*

Proof. By Lemma 3.3, \mathfrak{A} has a monster elementary extension \mathfrak{M} . Since \mathfrak{M} and \mathfrak{A} satisfy the same primitive positive sentences, $\text{CSP}(\mathfrak{A})$ is first-order definable if and only if $\text{CSP}(\mathfrak{M})$ is. The statement follows immediately from Theorem 5.7. \square

5.3. Horn definability. We will briefly examine a class of structures for which we can give a neat algebraic condition as to whether a relation that is qf-definable admits a qf Horn

definition. Recalling known complexity results for fo-expansions of $(\mathbb{R}; +, 1)$ we will see that the presence of a certain polymorphism exactly delineates those fo-expansions for which the CSP is NP-complete from those which are in P. The following proposition is essentially from [4].

Proposition 5.9. *Let \mathfrak{A} be a structure with an embedding e from \mathfrak{A}^2 into \mathfrak{A} . Then a relation R that is qf-definable in the relations of \mathfrak{A} is preserved by e iff it admits a qf-Horn definition in \mathfrak{A} .*

Proof. In this proof $\bar{x}_1 \dots, \bar{x}_k$ should be read as variable subtuples of the variable tuple \bar{x} . Likewise with the element subtuples $\bar{a}_1 \dots, \bar{a}_k$ of \bar{a} .

(Backwards.) Let \mathcal{F} be a Horn definition of R . Suppose \bar{a} and $\bar{a}' \in R^{\mathfrak{A}}$. It suffices to demonstrate the preservation of each clause in \mathcal{F} of the form $(R_1(\bar{x}_1) \wedge \dots \wedge R_l(\bar{x}_l)) \rightarrow R_{l+1}(\bar{x}_{l+1})$, for R_1, \dots, R_{l+1} relations of \mathfrak{A} .

$$\frac{\begin{array}{ccccccc} (R_1(\bar{a}_1) & \wedge & \dots & \wedge & R_l(\bar{a}_l) & \rightarrow & R_{l+1}(\bar{a}_{l+1}) \\ (R_1(\bar{a}'_1) & \wedge & \dots & \wedge & R_l(\bar{a}'_l) & \rightarrow & R_{l+1}(\bar{a}'_{l+1}) \end{array}}{(R_1(e(\bar{a}_1, \bar{a}'_1)) \wedge \dots \wedge R_l(e(\bar{a}_l, \bar{a}'_l))) \rightarrow R_{l+1}(e(\bar{a}_{l+1}, \bar{a}'_{l+1}))}$$

If the former clauses are true, there are two cases. Either some antecedent $R_i(\bar{a}_i)$ or $R_i(\bar{a}'_i)$ is false, in which case $R_i(e(\bar{a}_i, \bar{a}'_i))$ is false, and the latter clause is true. Or, if all antecedents in both former clauses are true, then both $R_{l+1}(\bar{a}_{l+1})$ and $R_{l+1}(\bar{a}'_{l+1})$ are true, so it follows that $R_{l+1}(e(\bar{a}_{l+1}, \bar{a}'_{l+1}))$ is true, and the latter clause is true.

(Forwards.) By contraposition. Consider a CNF definition \mathcal{F} of R in \mathfrak{A} that is irreducible in the sense that it has no redundant literals in its clauses. If it is not Horn, there exists a clause $R_1(\bar{x}_1) \vee R_2(\bar{x}_2) \vee S_3(\bar{x}_3) \vee \dots \vee S_l(\bar{x}_l)$, with R_1, R_2 positive literals S_3, \dots, S_l positive or negative literals, with $\bar{a}, \bar{a}' \in R^{\mathfrak{A}}$ s.t.

$$\begin{array}{l} R_1(\bar{a}_1) \wedge \neg R_2(\bar{a}_2) \wedge \neg S_3(\bar{a}_3) \wedge \dots \wedge \neg S_l(\bar{a}_l) \\ \neg R_1(\bar{a}'_1) \wedge R_2(\bar{a}'_2) \wedge \neg S_3(\bar{a}'_3) \wedge \dots \wedge \neg S_l(\bar{a}'_l) \end{array}$$

Consider the tuple $e(\bar{a}, \bar{a}')$. Clearly it will fail to satisfy the clause. □

We have already met an example of a structure with such an embedding: $(\mathbb{R}; +, 1)$.

Corollary 5.10. *Let \mathfrak{B} be an fo-expansion of $(\mathbb{R}; +, 1)$ and let $e : (\mathbb{R}; +, 1)^2 \rightarrow (\mathbb{R}; +, 1)$ be an embedding. Then: if e is a polymorphism of \mathfrak{B} , then $\text{CSP}(\mathfrak{B})$ is in P; otherwise $\text{CSP}(\mathfrak{B})$ is NP-complete.*

Proof. Note that $(\mathbb{R}; +, 1)$ admits quantifier elimination and so all fo-expansions may be specified as qf CNFs. It is proved in [8] that those that admit qf-Horn definitions give a CSP that is in P while those that do not give CSPs that are NP-complete. The result follows from Proposition 5.9. □

6. CONCLUDING REMARKS AND OPEN PROBLEMS

The universal-algebraic approach to the complexity of CSPs relies on two basic facts: the fact that every finite or ω -categorical structure is homomorphically equivalent to a model-complete core structure, and that primitive positive definability is characterized by preservation under polymorphisms.

In this paper we have presented generalizations of those two facts to not necessarily ω -categorical structures. A key concept in our proofs is the concept of existential-positive

closure, which we used to re-derive and generalize results in the literature about the presence and uniqueness of model-complete cores [3]. We also give a necessary and sufficient condition as to when a CSP can be formulated with an ω -categorical template.

The results of this paper show that the second part of the universal algebraic approach – the characterization of pp-definability by polymorphisms – can be applied to study the complexity of $\text{CSP}(\mathfrak{A})$ for arbitrary (and not just ω -categorical) infinite-domain structures \mathfrak{A} . Among one of the first applications, we

- gave a polymorphism-based characterization of those CSPs that are fo-definable;
- demonstrated in the context of real-valued constraint satisfaction problems that for large classes of CSPs the border between easy and hard constraint satisfaction can be described in terms of the existence of a certain polymorphism of the constraint language;
- have presented a strong universal-algebraic hardness criterion based on the absence of essential polymorphisms from the constraint language.

The following question is left for future research: can we strengthen our preservation theorem (Theorem 3.5) to show, under the additional assumption that \mathfrak{A} is epc (or epc and saturated, or has a model-complete core theory), that a first-order definable relation is pp-definable if and only if it is preserved by the *finitary* polymorphisms of \mathfrak{A} ?

ACKNOWLEDGEMENTS

We are grateful to our referees. Moreover, we would like to thank Moritz Müller for numerous comments which helped to improve the paper.

REFERENCES

- [1] BEN YAACOV, I. Positive model theory and compact abstract theories. *Journal of Mathematical Logic* 3, 1 (2003), 85–118.
- [2] BODIRSKY, M. Constraint satisfaction with infinite domains. Dissertation an der Humboldt-Universität zu Berlin, 2004.
- [3] BODIRSKY, M. Cores of countably categorical structures. *Logical Methods in Computer Science* 3, 1 (2007), 1–16.
- [4] BODIRSKY, M., CHEN, H., KÁRA, J., AND VON OERTZEN, T. Maximal infinite-valued constraint languages. *Theoretical Computer Science (TCS)* 410 (2009), 1684–1693. A preliminary version appeared at ICALP’07.
- [5] BODIRSKY, M., AND GROHE, M. Non-dichotomies in constraint satisfaction complexity. In *Proceedings of ICALP’08* (2008), pp. 184–196.
- [6] BODIRSKY, M., HERMANN, M., AND RICHOUX, F. Complexity of existential positive first-order logic. In *Proceedings of Computing in Europe (CiE’09)* (2009), pp. 31–36.
- [7] BODIRSKY, M., HILS, M., AND MARTIN, B. On the scope of the universal-algebraic approach to constraint satisfaction. In *Proceedings of LICS’10* (2010).
- [8] BODIRSKY, M., JONSSON, P., AND VON OERTZEN, T. Horn versus full first-order: a complexity dichotomy for algebraic constraint satisfaction problems. *To appear in the Journal of Logic and Computation* (2010).
- [9] BODIRSKY, M., AND KÁRA, J. The complexity of equality constraint languages. *Theory of Computing Systems* 3, 2 (2008), 136–158. A conference version appeared in the proceedings of CSR’06.
- [10] BODIRSKY, M., AND KÁRA, J. The complexity of temporal constraint satisfaction problems. In *Proceedings of the Symposium on Theory of Computing (STOC)* (2008), pp. 29–38.
- [11] BODIRSKY, M., AND NEŠETRIL, J. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation* 16, 3 (2006), 359–373.
- [12] BODNARČUK, V. G., KALUŽNIN, L. A., KOTOV, V. N., AND ROMOV, B. A. Galois theory for post algebras, part I and II. *Cybernetics* 5 (1969), 243–539.

- [13] BULATOV, A. Tractable conservative constraint satisfaction problems. In *Proceedings of LICS'03* (2003), pp. 321–330.
- [14] BULATOV, A. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *Journal of the ACM* 53, 1 (2006), 66–120.
- [15] BULATOV, A., KROKHIN, A., AND JEAVONS, P. G. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing* 34 (2005), 720–742.
- [16] BULATOV, A. A., AND VALERIOTE, M. Recent results on the algebraic approach to the CSP. In *Complexity of Constraints* (2008), pp. 68–92.
- [17] CREIGNOU, N., KOLAITIS, P. G., AND VOLLMER, H., Eds. *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar]* (2008), vol. 5250 of *Lecture Notes in Computer Science*, Springer.
- [18] FEDER, T., AND VARDI, M. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing* 28 (1999), 57–104.
- [19] GEIGER, D. Closed systems of functions and predicates. *Pacific Journal of Mathematics* 27 (1968), 95–100.
- [20] HOBBY, D., AND MCKENZIE, R. *The Structure of Finite Algebras*. AMS, 1988.
- [21] HODGES, W. *Model theory*. Cambridge University Press, 1993.
- [22] HODGES, W. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.
- [23] JEAVONS, P., COHEN, D., AND GYSENS, M. Closure properties of constraints. *Journal of the ACM* 44, 4 (1997), 527–548.
- [24] JEAVONS, P. G. On the algebraic structure of combinatorial problems. *Theoretical Computer Science* 200 (1998), 185–204.
- [25] LADNER, R. E. On the structure of polynomial time reducibility. *Journal of the ACM* 22, 1 (1975), 155–171.
- [26] LAROSE, B., LOTEN, C., AND TARDIF, C. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, DOI: 10.2168/LMCS-3(4:6) (2007).
- [27] POIZAT, B. Theorie de Galois pour les algebres de Post infinitaires. *Z. Math. Logik Grudl. Math.* 27 (1981), 31–44.
- [28] REINGOLD, O. Undirected connectivity in log-space. *Journal of the ACM* 55, 4 (2008).
- [29] ROMOV, B. Positive primitive structures. In *Multiple-Valued Logic, 2009. ISMVL '09. 39th International Symposium on* (may 2009), pp. 72–76.
- [30] ROMOV, B. A. Some properties of local partial clones on an infinite set. In *ISMVL '04: Proceedings of the 34th International Symposium on Multiple-Valued Logic* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 120–125.
- [31] ROSSMAN, B. Homomorphism preservation theorems. *J. ACM* 55, 3 (2008).
- [32] SIMMONS, H. Existentially closed structures. *J. Symb. Log.* 37, 2 (1972), 293–310.
- [33] SIMMONS, H. Large and small existentially closed structures. *Journal of Symbolic Logic* 41, 2 (1976), 379–390.
- [34] SZABÓ, L. Concrete representation of relational structures of universal algebras. *Acta Sci. Math. (Szeged)* 40 (1978), 175–184.
- [35] SZENDREI, A. *Clones in universal Algebra*. Séminaire de Mathématiques Supérieures. Les Presses de L'Université de Montréal, 1986.
- [36] ZILBER, B. *Uncountable categorical theories, Tranlations of Mathematical Monographs*, vol. 117. Amer. Math. Soc., 1993.