

Three Fundamental Questions in Modern Infinite-Domain Constraint Satisfaction

Michael Pinsker  

Technische Universität Wien, Austria

Jakub Rydval  

Technische Universität Wien, Austria

Moritz Schöbi  

Technische Universität Wien, Austria

Christoph Spiess  

Technische Universität Wien, Austria

Abstract

The Feder-Vardi dichotomy conjecture for Constraint Satisfaction Problems (CSPs) with finite templates, confirmed independently by Bulatov and Zhuk, has an extension to certain well-behaved infinite templates due to Bodirsky and Pinsker which remains wide open. We provide answers to three fundamental questions on the scope of the Bodirsky-Pinsker conjecture.

Our first two main results provide two simplifications of this scope, one of structural, and the other one of algebraic nature. The former simplification implies that the conjecture is equivalent to its restriction to templates without algebraicity, a crucial assumption in the most powerful classification methods. The latter yields that the higher-arity invariants of any template within its scope can be assumed to be essentially injective, and any algebraic condition characterizing any complexity class within the conjecture closed under Datalog reductions must be satisfiable by injections, thus lifting the mystery of the better applicability of certain conditions over others.

Our third main result uses the first one to show that any non-trivially tractable template within the scope serves, up to a Datalog-computable modification of it, as the witness of the tractability of a non-finitely tractable finite-domain Promise Constraint Satisfaction Problem (PCSP) by the so-called sandwich method. This generalizes a recent result of Mottet and provides a strong hitherto unknown connection between the Bodirsky-Pinsker conjecture and finite-domain PCSPs.

2012 ACM Subject Classification Theory of computation → Logic; Theory of computation → Computational complexity and cryptography

Keywords and phrases (Promise) Constraint Satisfaction Problem, dichotomy conjecture, polymorphism, identity, algebraicity, homogeneity, ω -categoricity, finite boundedness, Datalog

Digital Object Identifier 10.4230/LIPIcs.MFCS.2025.83

Related Version *Full Version:* <https://arxiv.org/abs/2502.06621>

Funding *Michael Pinsker, Jakub Rydval, and Moritz Schöbi:* This research was funded in whole or in part by the Austrian Science Fund (FWF) [I 5948]. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission. *Michael Pinsker and Christoph Spiess:* This research is funded by the European Union (ERC, POCOCOP, 101071674). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.



© Michael Pinsker, Jakub Rydval, Moritz Schöbi, and Christoph Spiess;
licensed under Creative Commons License CC-BY 4.0

50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025).

Editors: Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak; Article No. 83; pp. 83:1–83:20



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

1.1 The finite-domain CSP dichotomy theorem

Fixed-template *Constraint Satisfaction Problems* are computational problems parametrized by structures \mathfrak{A} with a finite relational signature, called *templates*; they are denoted by $\text{CSP}(\mathfrak{A})$ and ask whether a given structure \mathfrak{J} with the same signature as \mathfrak{A} admits a homomorphism to \mathfrak{A} . The general CSP framework is incredibly rich, in fact, it contains all decision problems up to polynomial-time equivalence [20]. For this reason, CSPs are typically studied under additional structural restrictions on the template \mathfrak{A} . The restriction that has received most attention is requiring the domain of \mathfrak{A} to be finite, and the problems arising in this way are called *finite-domain* CSPs. Already in this setting, one obtains many well-known problems such as HORN-SAT, 2-SAT, 3-SAT, or 3-COLORING.

Consider the example of 3-COLORING. This NP-complete problem can be modeled as $\text{CSP}(\mathfrak{K}_3)$, where \mathfrak{K}_3 stands for the complete graph on three vertices: indeed, a homomorphism to \mathfrak{K}_3 from a given structure \mathfrak{J} in the same signature (i.e. from a graph \mathfrak{J}) is simply a mapping that does not map any two adjacent vertices of \mathfrak{J} to the same vertex of \mathfrak{K}_3 . A key observation in the theory of CSPs is that the NP-completeness of $\text{CSP}(\mathfrak{K}_3)$ can be extended to CSPs of other structures which can “simulate” \mathfrak{K}_3 over their domain or over a finite power thereof using *primitive positive* (pp) formulas. Enhancing this by the additional observation that *homomorphically equivalent* templates have equal CSPs, one arrives at the notion of *pp-constructibility* [9]. If \mathfrak{A}' is pp-constructible from \mathfrak{A} , then $\text{CSP}(\mathfrak{A}')$ is reducible to $\text{CSP}(\mathfrak{A})$ in logarithmic space; moreover, this reduction can be formulated in the logic programming language *Datalog*. This uniform reduction between CSP templates is sufficiently powerful, as we will see below, to describe all NP-hardness amongst finite-domain CSPs.

In the early 2000s, the field of finite-domain constraint satisfaction quickly rose in fame, in particular due to the discovery of tight connections with universal algebra. The most basic of these connections is that two structures with identical domains have the same sets of pp-definable relations if and only if they have the same sets of polymorphisms [52, 51]. Here, a *polymorphism* of a relational structure \mathfrak{A} is simply a homomorphism from a finite power of \mathfrak{A} into \mathfrak{A} itself; we denote by $\text{Pol}(\mathfrak{A})$ the set of all polymorphisms of \mathfrak{A} , the *polymorphism clone* of \mathfrak{A} . Since taking expansions by pp-definable relations does not lead to an increase in complexity, the study of finite-domain CSPs is subsumed by the study of finite algebras. Over the past two decades, the link between universal algebra and CSPs was gradually refined. After an intermediate stop at *pp-interpretations* [37], today we know that also pp-constructibility between finite relational structures is fully encoded in their polymorphisms [9, Theorem 1.3]: \mathfrak{A}' is pp-constructible from \mathfrak{A} if and only if $\text{Pol}(\mathfrak{A}')$ satisfies all *height-1 identities* of $\text{Pol}(\mathfrak{A})$. By a breakthrough result of Bulatov [36] and Zhuk [68, 69], the pp-construction of \mathfrak{K}_3 is the unique source of NP-hardness for finite-domain CSPs (if $\text{P} \neq \text{NP}$), and in fact, they provided a polynomial-time algorithm for any such CSP which does not pp-construct \mathfrak{K}_3 . Besides its complexity-theoretic significance (in particular, of confirming the dichotomy conjecture of Feder and Vardi [44]), this result has the additional appeal that the border for tractability can be described by a neat universal-algebraic condition on polymorphism clones. We state the theorem of Bulatov and Zhuk in a formulation which takes into account the results in [9] and [66].

► **Theorem 1** (Bulatov and Zhuk [36, 68]). *Let \mathfrak{A} be a finite relational structure. Then either \mathfrak{A} pp-constructs \mathfrak{K}_3 and $\text{CSP}(\mathfrak{A})$ is NP-complete, or $\text{Pol}(\mathfrak{A})$ satisfies the Siggers identity $s(x, y, z, x, y, z) \approx s(y, z, x, z, x, y)$ and $\text{CSP}(\mathfrak{A})$ is solvable in polynomial time.*

1.2 The infinite-domain CSP tractability conjecture

Arguably, the two main reasons for the popularity of finite-domain CSPs over CSPs which require an infinite template are immediate containment of such problems in the class NP (a homomorphism can be guessed and verified in polynomial time) as well as the above-mentioned applicability of algebraic methods which had been developed independently of CSPs for decades. Yet, even within the class NP, finite-domain CSPs are of an extremely restricted kind: already simple problems such as ACYCLICITY of directed graphs (captured by the template $(\mathbb{Q}; <)$), various natural coloring problems for graphs such as NO-MONOCHROMATIC-TRIANGLE, and more generally the model-checking problem for natural restrictions of existential second-order logic such as the logic MMSNP of Feder and Vardi [44], are beyond its primitive scope despite their containment in NP.

The quest for a CSP framework including such problems whilst staying within NP and allowing for an algebraic approach akin to the one for finite templates started with the popularization of ω -categoricity by Bodirsky and Nešetřil [27] as a sufficient structural restriction ensuring the latter: for countable ω -categorical structures, pp-definability is determined by their polymorphisms and, as was subsequently shown, so are the general reduction of pp-interpretability [29] as well as the pp-constructibility of \mathfrak{K}_3 [9]. The ω -categoricity of \mathfrak{A} can be interpreted as \mathfrak{A} being “finite modulo automorphisms”: more precisely, the action of its automorphism group $\text{Aut}(\mathfrak{A})$ on d -tuples has only finitely many orbits for all finite $d \geq 1$. It is, however, a mathematical property with little computational bearing: the CSPs of ω -categorical structures can be monstrous from this perspective [49, 48] (complete for a variety of complexity classes of arbitrarily high complexity). For this reason, and based on strong empirical evidence, Bodirsky and Pinsker [31] identified a proper subclass of countable ω -categorical structures as a candidate for an algebraic complexity dichotomy extending the theorem of Bulatov and Zhuk which does not seem to suffer from this deficiency.

Their first requirement on \mathfrak{A} is that the orbit under $\text{Aut}(\mathfrak{A})$ of any d -tuple is determined by the relations that hold on it: we call a relational structure \mathfrak{A} *homogeneous* if every isomorphism between its finite substructures extends to an automorphism of \mathfrak{A} . This gives an effective way of representing orbits, which are central to ω -categoricity. They then additionally impose an effective way of determining whether a given finite substructure is contained in \mathfrak{A} , which places the CSP in NP: a relational structure \mathfrak{A} over a finite signature τ is *finitely bounded* [55] if there exists a finite set \mathcal{N} of finite τ -structures (*bounds*) such that a finite τ -structure embeds into \mathfrak{A} if and only if it does not embed any member of \mathcal{N} ; equivalently, the finite substructures are up to isomorphism precisely the finite models of some universal first-order sentence [14, Lemma 2.3.14] (see also [25, 64] for more details).

► **Example 2.** A standard example of a countable finitely bounded homogeneous structure is $(\mathbb{Q}; <)$. Its finite substructures are all finite strict linear orders, which can be axiomatized by irreflexivity, totality, and transitivity: $\forall x, y, z (\neg(x < x) \wedge (x < y \vee x = y \vee y < x) \wedge (x < y \wedge y < z \Rightarrow x < z))$. Moreover, every isomorphism between two such finite substructures can be extended to an automorphism of $(\mathbb{Q}; <)$, e.g. by a piecewise affine transformation.

Finite boundedness and homogeneity taken together is a very strong assumption. Bodirsky and Pinsker observed that for all practical purposes, in particular the applicability of polymorphisms to complexity as well as containment in NP, it is sufficient that the template \mathfrak{A} be a *first-order reduct* of a structure \mathfrak{B} enjoying these, i.e. first-order definable therein. This yields sufficient flexibility to model a huge class of computational problems which includes in particular the ones mentioned above. In fact, requiring \mathfrak{A} to be a *reduct* of a finitely bounded homogeneous structure \mathfrak{B} , i.e. obtained from \mathfrak{B} by forgetting relations, turns out to be equivalent [3, Proposition 7], and we shall find this approach convenient.

► **Example 3.** The CSP of the template $(\mathbb{Q}; \{(x, y, z) \mid x < y < z \vee z < y < x\})$, first-order definable in $(\mathbb{Q}; <)$, is the classical NP-complete *betweenness problem* from the 1970s [61].

The following formulation of the conjecture takes into account later progress [9, 10, 11, 7].

► **Conjecture 4** (Bodirsky and Pinsker 2012 [31]). *Let \mathfrak{A} be a reduct of a countable finitely bounded homogeneous structure \mathfrak{B} . Then exactly one of the following holds.*

- \mathfrak{A} pp-constructs \mathfrak{K}_3 (and consequently, $\text{CSP}(\mathfrak{A})$ is NP-complete);
- \mathfrak{A} does not pp-construct \mathfrak{K}_3 ; in this case $\text{Pol}(\mathfrak{A})$ satisfies the pseudo-Siggers identity $\alpha \circ s(x, y, z, x, y, z) \approx \beta \circ s(y, z, x, z, x, y)$ and $\text{CSP}(\mathfrak{A})$ is polynomial-time solvable.

It is important to note that the open part of Conjecture 4 is only the consequence of polynomial-time tractability: the negation of the first item does yield the satisfaction of the pseudo-Siggers identity by a theorem due to Barto and Pinsker [11]. It is even equivalent to it for *model-complete cores* [7, Theorem 1.3]: an ω -categorical model-complete core is a template \mathfrak{A} whose endomorphisms preserve all orbits of $\text{Aut}(\mathfrak{A})$; it exists for all CSPs with an ω -categorical template and is unique up to isomorphism [13].

2 Three fundamental questions: our contributions

The present paper aims to answer three fundamental questions around Conjecture 4, of which the first concerns its scope; the second the algebraic invariants of templates therein; and the third its connection with the rapidly evolving field of (finite-domain) Promise Constraint Satisfaction Problems. More precisely, we consider the following questions:

1. Are there significant additional structural assumptions, perhaps of model-theoretic nature, that can be imposed onto the structures of Conjecture 4 without loss of generality, i.e. without affecting the truth of the conjecture?
2. Are there significant algebraic assumptions that can be imposed onto the polymorphisms of the structures of Conjecture 4 without loss of generality?
3. Are there algorithmic connections between CSPs from Conjecture 4 and finite-domain Promise Constraint Satisfaction Problems?

We provide affirmative answers to all three questions, as we shall describe in the following.

2.1 Algebraicity is irrelevant

Conjecture 4 has been confirmed for many subclasses. Mottet and Pinsker speak of first- and second-generation classifications [60]: those of the former kind can roughly be described as exhaustive case-by-case analyzes, using Ramsey theory, of the available polymorphisms for the first-order reducts of a fixed finitely bounded homogeneous structure \mathfrak{B} – extensive complexity classifications such as those for temporal CSPs [22], Graph-SAT problems [28], Poset-SAT problems [54], and CSPs with templates first-order definable in arbitrary homogeneous graphs [24] were obtained this way. Second-generation classifications take a more structured approach mimicking advanced algebraic methods for finite domains; they were employed, for example, to achieve classifications for the logic MMSNP [23], Hypergraph-SAT problems [59], and certain graph orientation problems [45, 12]. However, even the most advanced methods as described in [60] require additional abstract structural assumptions on the template.

One such assumption is that the template \mathfrak{A} , or even the structure \mathfrak{B} in which it is first-order defined, has *no algebraicity*. It is present virtually everywhere in the infinite-domain CSP literature [23]; and in particular an important prerequisite in most of the general results of the theory of *smooth approximations* developed by Mottet and Pinsker [60], as well as in

all results of Bodirsky and Greiner [18, 19] about CSPs of *generic superpositions* of theories. We say that a structure \mathfrak{A} has *no algebraicity* (in the group-theoretic sense) if, for every $k \geq 1$ and every $\bar{a} \in A^k$, the automorphisms of \mathfrak{A} which stabilize \bar{a} do not stabilize any $a' \notin \bar{a}$ (i.e. a' does not appear as an entry of \bar{a}). In ω -categorical structures this is the case if and only if, for every $k \geq 1$, every tuple $\bar{a} \in A^k$, and every $a' \notin \bar{a}$, it is not possible to first-order define the unary relation $\{a'\}$ using \bar{a} as parameters; in other words, any non-trivial first-order property relative to \bar{a} is always satisfied by infinitely many elements, which implies that some arguments become easier compared to finite structures (whereas naturally, many other arguments become harder). We will show that the following strengthening of this property (for model-complete cores, see Proposition 13) can be assumed when resolving Conjecture 4: a structure \mathfrak{A} is *CSP-injective* if every finite structure that homomorphically maps to \mathfrak{A} also does so injectively; in other words, if there is a solution to an instance of $\text{CSP}(\mathfrak{A})$, then there is also an injective one. CSP-injectivity played an essential role in the universal-algebraic proof of the complexity dichotomy for Monotone Monadic SNP [23] and, more recently, certain finitely bounded homogeneous uniform hypergraphs [59].

► **Example 5.** The template $(\mathbb{Q}; <)$ from Example 2 is CSP-injective: a digraph maps homomorphically to $(\mathbb{Q}; <)$ if and only if it is acyclic; in this case there is always an injective homomorphism. It is a model-complete core, and hence has no algebraicity (Proposition 13). Its expansion $(\mathbb{Q}; <)_{I_4}$ by the relation $I_4 := \{(x, y, u, v) \mid x = y \Rightarrow u = v\}$ has the same automorphisms, and hence no algebraicity; it is no longer CSP-injective (Proposition 14). The equivalence relation with countably many classes of size 2 has algebraicity: any automorphism stabilizing some element a must also stabilize the other element of its class.

One of the properties preserved by our construction enforcing CSP-injectivity in Theorem 6 below is that of the structure \mathfrak{B} being *Ramsey*. Although of central importance in classification methods [30, 62], its precise definition is not essential to the present paper and therefore omitted; we refer the reader to [50] for details about structural Ramsey theory.

Throughout the present paper, the notation $\mathfrak{A} \wr \mathfrak{B}$ stands for the relational wreath product of two structures \mathfrak{A} and \mathfrak{B} , introduced formally in Section 4.1, and \mathfrak{A}_R stands for the expansion of a structure \mathfrak{A} by a relation R defined over its domain.

► **Theorem 6 (No Algebraicity).** *Let \mathfrak{A} be a non-trivial reduct of a countable structure \mathfrak{B} over a finite relational signature. Then $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ is a non-trivial CSP-injective reduct of $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$, which is a countably infinite structure without algebraicity over a finite relational signature, and $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq})$ and $\text{CSP}(\mathfrak{A})$ are Datalog-interreducible. Moreover:*

1. *If \mathfrak{B} is ω -categorical, then $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ is ω -categorical as well. In this case, $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ is a model-complete core if and only if \mathfrak{A} is a model-complete core.*
2. *If \mathfrak{B} is homogeneous, then $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ is homogeneous as well. In this case, $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ is Ramsey if and only if \mathfrak{B} is Ramsey.*
3. *If \mathfrak{B} is finitely bounded, then $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ is finitely bounded as well.*
4. *If the number of orbits of d -tuples of \mathfrak{A} has growth slower than 2^{2^d} , then $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ pp-constructs \mathfrak{K}_3 if and only if \mathfrak{A} pp-constructs \mathfrak{K}_3 .*

It is not hard to see that item 4 of Theorem 6 applies to reducts of homogeneous structures in a finite relational signature [55], in particular to all structures in the scope of Conjecture 4.

2.2 Polymorphisms are injective

Comparing Theorem 1 with Conjecture 4, one notices the replacement of the Siggers identity by its (weaker) pseudo-variant. This is not just an inconvenience arising from the proof of Barto and Pinsker in [11] (or the more recent proof in [5]), but there are examples showing the necessity to weaken the condition.

► **Example 7.** The CSP of $(\mathbb{Q}; <)_{I_4}$ from Example 5 is solvable in Datalog (Sections 3 and 4.2). Since $\text{CSP}(\mathfrak{K}_3)$ is not, and pp-constructions provide Datalog-reductions [1], it follows that $(\mathbb{Q}; <)_{I_4}$ does not pp-construct \mathfrak{K}_3 . Hence its polymorphisms satisfy the pseudo-Siggers identity by the above-mentioned theorem of Barto and Pinsker. However, any polymorphism of I_4 is injective up to dummy variables [15], and hence cannot satisfy the Siggers identity.

We say that a template is *Pol-injective* if all of its polymorphisms are *essentially injective*, i.e. injective up to dummy variables – this is the case if and only if the relation I_4 is invariant under its polymorphisms. As it turns out, the relation I_4 can be added to the templates constructed in Theorem 6 at no computational cost (up to Datalog reductions). It thus follows that any algebraic condition on polymorphisms capturing a complexity class closed under Datalog-reductions for CSPs of structures within the scope of Conjecture 4 must necessarily be satisfiable by essentially injective operations. Until now, various concrete examples of natural templates all of whose polymorphisms are essentially injective had pointed towards a statement of this kind (see e.g. [7]); we provide a rigorous confirmation.

► **Theorem 8.** *The structures $(\mathbb{Q}; \wr \mathfrak{A})_{\neq, I_4}$ and $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq, I_4}$ enjoy all properties of $(\mathbb{Q}; \wr \mathfrak{A})_{\neq}$ and $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ from Theorem 6 except that CSP-injectivity is replaced by Pol-injectivity.*

For finite structures the only essentially injective polymorphisms are at most unary up to dummy variables. It is therefore of no surprise that so few universal-algebraic conditions for CSPs from the finite have been successfully lifted to the infinite. One prominent example which can actually be lifted is the satisfiability of *quasi near-unanimity identities*, which captures *bounded strict width* both in the finite and in the infinite [16]; bounded strict width corresponds to solvability in a particular proper fragment of Datalog that is not closed under the Datalog-reductions in Theorem 6. Examples of conditions for polymorphism clones that can be satisfied by essentially injective operations are any *pseudo height-1 identities* (e.g., the pseudo-Siggers identity), but also the *dissected near-unanimity identities* [48] which were proposed in [32] as a possible candidate for solvability of CSPs in *fixed-point logic*.

Marimon and Pinsker have recently provided a negative answer, within the scope of Conjecture 4, to a question of Bodirsky [14, Question 14.2.6(27)] asking whether every ω -categorical CSP template without algebraicity which is solvable in Datalog has a binary injective polymorphism [56, Corollary 8.10]. We contrast this negative result with a corollary to Theorem 8 that implies that up to Datalog reductions the answer is positive.

► **Corollary 9.** *Let \mathfrak{A} be a non-trivial reduct of a finitely bounded homogeneous structure \mathfrak{B} . If \mathfrak{A} does not pp-construct \mathfrak{K}_3 , then $(\mathbb{Q}; \wr \mathfrak{A})_{\neq, I_4}$ has a binary injective polymorphism.*

2.3 Everything is a cheese

One of the currently most vibrant branches of research in constraint satisfaction are *Promise Constraint Satisfaction Problems* (PCSPs). For two relational structures \mathfrak{S}_1 and \mathfrak{S}_2 such that \mathfrak{S}_1 maps homomorphically to \mathfrak{S}_2 , the problem $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ asks to decide whether a given finite structure \mathfrak{J} in the same signature as \mathfrak{S}_1 and \mathfrak{S}_2 maps homomorphically to \mathfrak{S}_1

or does not even map homomorphically to \mathfrak{S}_2 ; the instance \mathfrak{J} is promised to satisfy one of those two cases. Thus, the problem $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ asks to compare $\text{CSP}(\mathfrak{S}_1)$ to a relaxation $\text{CSP}(\mathfrak{S}_2)$ of it, and since $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_1) = \text{CSP}(\mathfrak{S}_1)$, promise constraint satisfaction problems generalize the CSP framework. This generalization is vast, and contains many well-known computational problems such as APPROXIMATE GRAPH COLORING [46], $(2 + \varepsilon)$ -SAT [2], and HYPERGRAPH COLORING [41]. Although formally not necessary, research has focused on finite templates $(\mathfrak{S}_1, \mathfrak{S}_2)$: this is justified for example by the fact that even for structures over a Boolean domain, there is no complete complexity classification yet. Hence, the two extensions considered here of finite-domain CSPs to ω -categorical templates, or alternatively to PCSPs, are somewhat orthogonal, and it is natural to wonder whether there exist any connections, in particular of algorithmic nature.

One potential link is the use of ω -categorical CSP templates as cheeses in the sandwich method to prove tractability of PCSPs, as follows. Note that if \mathfrak{A} is a structure such that \mathfrak{S}_1 maps homomorphically to \mathfrak{A} and \mathfrak{A} maps homomorphically to \mathfrak{S}_2 , then $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ trivially reduces to $\text{CSP}(\mathfrak{A})$; we call \mathfrak{A} a *cheese sandwiched* by the template $(\mathfrak{S}_1, \mathfrak{S}_2)$. This simple observation is more powerful than it might first appear: in fact, so far, every known tractable (finite-domain) PCSP can be reduced to the tractable CSP of some cheese structure via the sandwich method [57, 40]. On the other hand, Barto [4] provided an example of a PCSP template sandwiching an infinite polynomial-time tractable cheese while not being *finitely tractable*, i.e. such that every finite-domain cheese has an NP-complete CSP [4]. Since Barto's cheese is not ω -categorical, he raised the question whether such an example could also be constructed with an ω -categorical cheese ([4, Question IV.1]). Subsequently the question arose, and was promoted in particular by Zhuk at the CSP World Congress 2023, whether structures within the scope of Conjecture 4 could *ever* serve as tractable cheeses in the absence of a finite one. The first two examples of tractable ω -categorical cheeses for non-finitely tractable finite-domain PCSPs were recently obtained by Mottet [58]; his structures fall within the scope of Conjecture 4. In the present paper, we strengthen this result by creating non-finitely tractable finite-domain PCSPs from virtually every structure \mathfrak{A} within the scope of Conjecture 4. Our construction consists of two steps. First, we show that such a construction is possible in general under the additional assumption that \mathfrak{A} is equipped with an inequality predicate and that it is a reduct of a structure \mathfrak{B} which is linearly ordered. We then combine this construction with Theorem 6. Modulo an extra step consisting of taking *generic superpositions* (Proposition 18) of structures with $(\mathbb{Q}; <)$ (for which we need the property of no algebraicity), this finishes our proof.

► **Theorem 10** (Datalog-Approximability of Infinite-Domain CSPs). *For every non-trivial τ -reduct \mathfrak{A} of a countable finitely bounded homogeneous structure \mathfrak{B} there exists a τ' -reduct \mathfrak{A}' of a countable finitely bounded homogeneous structure \mathfrak{B}' together with Datalog-interpretations \mathcal{J} and \mathcal{J}' from τ -structures to τ' -structures and vice versa such that, for every $n \in \mathbb{N}$, there exists a finite PCSP template $(\mathfrak{S}_1, \mathfrak{S}_2)$ with the following properties:*

1. \mathfrak{A} pp-constructs \mathfrak{R}_3 if and only if \mathfrak{A}' pp-constructs \mathfrak{R}_3 .
2. \mathcal{J} and \mathcal{J}' are (Datalog-)reductions from $\text{CSP}(\mathfrak{A})$ to $\text{CSP}(\mathfrak{A}')$ and vice versa.
3. If \mathfrak{A} is a model-complete core, then \mathfrak{A}' is a model-complete core as well.
4. If \mathfrak{B} is Ramsey, then \mathfrak{B}' is Ramsey as well.
5. Every finite cheese for $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ pp-constructs \mathfrak{R}_3 ; and hence $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ is not finitely tractable.
6. \mathfrak{A}' is a cheese for $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$; and hence $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ reduces to $\text{CSP}(\mathfrak{A})$ via \mathcal{J}' .
7. Conversely, \mathcal{J} is correct as a reduction from $\text{CSP}(\mathfrak{A})$ to $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ when limited to instances \mathfrak{X} of $\text{CSP}(\mathfrak{A})$ with the property that \mathfrak{X} homomorphically maps to \mathfrak{A} if and only if \mathfrak{X} homomorphically maps to a substructure of \mathfrak{A} of domain size $\leq n$.

The last item in Theorem 10 is not necessary for obtaining non-finitely tractable PCSPs from Conjecture 4, but provides an additional link that makes it possible to transfer logical inexpressibility results from infinite-domain CSPs to PCSPs. For example, the PCSPs provided by the theorem from $\text{CSP}(\mathbb{Q}; X)$, where $X = \{(x, y, z) \mid x = y < z \vee y = z < x \vee z = x < y\}$, will not be solvable in *Fixed-Point logic with Counting* (FPC), because $\text{CSP}(\mathbb{Q}; X)$ is not [32] and the instances witnessing the inexpressibility in FPC homomorphically map to $(\mathbb{Q}; X)$ if and only if they homomorphically map to a fixed finite substructure of it (cf. [58, Proposition 37]). We expect a similar situation for the PCSPs associated with the first-order reducts of the *homogeneous C-relation* [21], due to the recent work [67, Theorem 4.7].

Our results compare to the above-mentioned examples of Mottet [58, Propositions 35 and 36] of a tractable ω -categorical cheese for a finite-domain PCSP that is not finitely tractable as follows. First, Mottet's finite-domain PCSPs [58, Theorem 2] are reducible to the associated ω -categorical templates under *gadget reductions*, which are the specific form of Datalog-reductions stemming from pp-constructions. The main advantage of gadget reductions over general Datalog-reductions is that, in many settings, gadget reductions between (P)CSPs are captured by the transfer of height-1 identities between polymorphism clones (or more generally *polymorphism minions*); see [39]. As a trade-off for the generality of our result (Theorem 10), our Datalog-reductions are not given in terms of gadgets; in fact, one can formally prove that, in many cases, our finite-domain PCSPs do not admit a gadget reduction to the original ω -categorical template (Example 15). Second, Mottet's work [58] also contains a result (Theorem 1) showing that every CSP within the scope of the Bodirsky-Pinsker conjecture is polynomial-time equivalent to the PCSP of a half-infinite template, more precisely a template $(\mathfrak{S}_1, \mathfrak{S}_2)$ such that \mathfrak{S}_1 is in the scope of the Bodirsky-Pinsker conjecture and \mathfrak{S}_2 is finite. In contrast, our Theorem 10 provides non-finitely tractable finite-domain PCSP templates $(\mathfrak{S}_1, \mathfrak{S}_2)$ sandwiching a given structure \mathfrak{A} from Conjecture 4 up to Datalog interreducibility. Since our PCSP templates $(\mathfrak{S}_1, \mathfrak{S}_2)$ are finite and not half-infinite, we cannot guarantee polynomial-time equivalence between $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ and $\text{CSP}(\mathfrak{A})$. However, our construction still allows us to approximate $\text{CSP}(\mathfrak{A})$ to an arbitrary degree in the sense of item 7 in Theorem 10.

3 Basic definitions

The set $\{1, \dots, n\}$ is denoted by $[n]$, and we use the bar notation \bar{t} for tuples. The *component-wise action* of $f: A^n \rightarrow B$ on k -tuples $(x_{1,1}, \dots, x_{1,k}), \dots, (x_{n,1}, \dots, x_{n,k}) \in A^k$ is the k -tuple $f((x_{1,1}, \dots, x_{1,k}), \dots, (x_{n,1}, \dots, x_{n,k})) := (f(x_{1,1}, \dots, x_{n,1}), \dots, f(x_{1,k}, \dots, x_{n,k})) \in B^k$.

Relational structures. A *(relational) signature* τ is a set of *relation symbols*, each $R \in \tau$ with an associated natural number called *arity*. A *(relational) τ -structure* \mathfrak{A} consists of a set A (the *domain*) together with the relations $R^{\mathfrak{A}} \subseteq A^k$ for each $R \in \tau$ with arity k . An *expansion* of \mathfrak{A} is a σ -structure \mathfrak{B} with $A = B$ such that $\tau \subseteq \sigma$ and $R^{\mathfrak{B}} = R^{\mathfrak{A}}$ for each relation symbol $R \in \tau$. Conversely, we call \mathfrak{A} a *reduct* of \mathfrak{B} and denote it by $\mathfrak{B}|^{\tau}$. For a positive integer d and a τ -structure \mathfrak{A} , the *d -th power* of \mathfrak{A} is the τ -structure \mathfrak{A}^d with domain A^d and relations $R^{\mathfrak{A}^d} = \{(a_{1,1}, \dots, a_{1,d}), \dots, (a_{k,1}, \dots, a_{k,d}) \in (A^d)^k \mid (a_{1,1}, \dots, a_{k,1}), \dots, (a_{1,d}, \dots, a_{k,d}) \in R^{\mathfrak{A}}\}$ for each $R \in \tau$ of arity k . The *substructure* of a τ -structure \mathfrak{A} on a subset $B \subseteq A$ is the τ -structure \mathfrak{B} with domain B and relations $R^{\mathfrak{B}} = R^{\mathfrak{A}} \cap B^k$ for every $R \in \tau$ of arity k . The *factor* of a τ -structure \mathfrak{A} through an equivalence relation $E \subseteq A^2$ is the τ -structure \mathfrak{A}/E with domain A/E and relations $R^{\mathfrak{A}/E} = q_E(R^{\mathfrak{A}})$, where q_E denotes the factor map $x \mapsto [x]_E$.

A *homomorphism* $h: \mathfrak{A} \rightarrow \mathfrak{B}$ for τ -structures \mathfrak{A} and \mathfrak{B} is a mapping $h: A \rightarrow B$ that *preserves* each τ -relation, i.e., if $\bar{t} \in R^{\mathfrak{A}}$ for some relation symbol $R \in \tau$, then $h(\bar{t}) \in R^{\mathfrak{B}}$. We write $\mathfrak{A} \rightarrow \mathfrak{B}$ if \mathfrak{A} maps homomorphically to \mathfrak{B} ; formally, $\text{CSP}(\mathfrak{A}) = \{\mathfrak{J} \text{ finite} \mid \mathfrak{J} \rightarrow \mathfrak{A}\}$. If $\mathfrak{A} \rightarrow \mathfrak{B}$ and $\mathfrak{B} \rightarrow \mathfrak{A}$, then we say that \mathfrak{A} and \mathfrak{B} are *homomorphically equivalent*. An *endomorphism* of \mathfrak{A} is a homomorphism from \mathfrak{A} to itself; we denote by $\text{End}(\mathfrak{A})$ the set (monoid) of all endomorphisms of \mathfrak{A} . Clearly, if \mathfrak{A} has a constant endomorphism, then its CSP is trivial, and we call \mathfrak{A} *trivial* itself; otherwise we call \mathfrak{A} *non-trivial*. As mentioned in the introduction, a *polymorphism* of \mathfrak{A} is a homomorphism from \mathfrak{A}^k into \mathfrak{A} for some $k \in \mathbb{N}$.

An *embedding* is an injective homomorphism $h: \mathfrak{A} \rightarrow \mathfrak{B}$ that additionally satisfies the following condition: for every k -ary relation symbol $R \in \tau$ and $\bar{t} \in A^k$ we have $h(\bar{t}) \in R^{\mathfrak{B}}$ only if $\bar{t} \in R^{\mathfrak{A}}$. An *isomorphism* is a surjective embedding. Two structures \mathfrak{A} and \mathfrak{B} are *isomorphic* if there exists an isomorphism from \mathfrak{A} to \mathfrak{B} . An *automorphism* is an isomorphism from \mathfrak{A} to \mathfrak{A} ; we denote by $\text{Aut}(\mathfrak{A})$ the set (group) of all automorphisms of \mathfrak{A} . The *orbit* of a tuple $\bar{t} \in A^k$ in \mathfrak{A} is the set $\{g(\bar{t}) \mid g \in \text{Aut}(\mathfrak{A})\}$. Any tuples belonging to the same orbit satisfy precisely the same first-order formulas over \mathfrak{A} . The *orbit equivalence* relation provides a natural way to factorise relational structures; in the present paper, we will only need the following specific case. Given a relational structure \mathfrak{A} and a subgroup $G \subseteq \text{Aut}(\mathfrak{A})$, we denote by $\mathfrak{A}/_G$ the structure \mathfrak{A}/E for $E := \{(x, y) \in A^2 \mid \exists g \in G: g(x) = y\}$.

Universal algebra. An (*equational*) *condition* is a set of *identities*, i.e. formal expressions of the form $s \approx t$, where s and t are terms over a common set of function symbols. An equational condition is *height-1* if it contains neither nested terms nor terms consisting of a single variable. We say that $\text{Pol}(\mathfrak{A})$ (or some set of operations on A) *satisfies* an equational condition Σ if the function symbols can be interpreted as elements of $\text{Pol}(\mathfrak{A})$ so that, for each identity $s \approx t$ in Σ , the equality $s = t$ holds for any evaluation of variables in A . An example is the *cyclic identity* of arity $n \geq 2$ given by $f(x_1, \dots, x_n) \approx f(x_2, \dots, x_n, x_1)$, for a symbol f of arity $n \geq 2$. We say that an operation is *cyclic* if it satisfies the cyclic identity.

► **Theorem 11** ([8]). *Let \mathfrak{A} be a finite relational structure. If \mathfrak{A} does not pp-construct \mathfrak{K}_3 , then $\text{Pol}(\mathfrak{A})$ contains a cyclic operation.*

A *pseudo-version* of a height-1 identity $s \approx t$ is of the form $\alpha \circ s \approx \beta \circ t$ for fresh unary function symbols α and β . An example is the pseudo-Siggers identity from Conjecture 4. This definition naturally extends to height-1 conditions.

We extend the notion of a polymorphism to PCSP templates $(\mathfrak{S}_1, \mathfrak{S}_2)$ in the usual way: a *polymorphism* of $(\mathfrak{S}_1, \mathfrak{S}_2)$ is a homomorphism from a finite power of \mathfrak{S}_1 to \mathfrak{S}_2 ; we denote by $\text{Pol}(\mathfrak{S}_1, \mathfrak{S}_2)$ the set (minion) of all such polymorphisms. Also the satisfaction of height-1 conditions can be generalized to the PCSP setting in a straightforward manner. In the case of nested identities, one has to be more careful because the composition of polymorphisms between different structures is not defined. We say that a height-1 condition Σ is *satisfied in $\text{Pol}(\mathfrak{S}_1, \mathfrak{S}_2)$ modulo a set H of unary operations on S_2* if the function symbols in Σ can be interpreted as elements of $\text{Pol}(\mathfrak{S}_1, \mathfrak{S}_2)$ such that for each identity $s \approx t$ in Σ there exist $\alpha, \beta \in H$ such that $\alpha \circ s = \beta \circ t$ holds for any evaluation of variables in S_1 .

Datalog. Datalog reductions between CSPs are typically specified using factor-free Datalog-interpretations with parameters [1, 39], which are a particular case of logical interpretations with parameters. In the present paper, we only need a very basic understanding of Datalog-reductions, which can be easily explained on an informal level.

Datalog is defined by adding formation rules to the existential positive fragment of first-order logic whose semantics is defined with inflationary fixed-points of definable operators. Every existential positive first-order formula is a Datalog formula and, if $\phi(\bar{x})$ is a Datalog formula over some relational signature $\tau \cup \{R\}$ with $R \notin \tau$, then $[\text{ifp}_R \phi](\bar{x})$ is a Datalog formula over the signature τ whose semantics is given as follows. For a τ -structure \mathfrak{A} and a tuple \bar{a} over A matching the arity of \bar{x} , say k , we have $\mathfrak{A} \models [\text{ifp}_R \phi](\bar{a})$ if and only if \bar{a} is contained in the inflationary fixed-point of the operator $F_{\phi, R}^{\mathfrak{A}}(X) := \{\bar{x} \in A^k \mid \mathfrak{A}_X \models \phi(\bar{x})\}$, i.e., the limit of the sequence $X_0 := \emptyset$ and $X_{i+1} := X_i \cup F_{\phi, R}^{\mathfrak{A}}(X_i)$. For example, the Datalog formula $[\text{ifp}_T (x < z) \vee (\exists y. T(x, y) \wedge T(y, z))](x, z)$ computes the transitive closure of $<$.

For a relational τ -structure \mathfrak{A} , we say that $\text{CSP}(\mathfrak{A})$ is *solvable in Datalog* if there exists a Datalog sentence defining the class of all finite τ -structures which do *not* homomorphically map to \mathfrak{A} . For example, $\text{CSP}(\mathbb{Q}; <)$ is solvable in Datalog using the sentence $\exists u[\text{ifp}_T (x < z) \vee (\exists y. T(x, y) \wedge T(y, z))](u, u)$. It is not hard to see that every Datalog formula is specified by a finite set of rules of the form $R(\bar{x}) \Leftarrow R_1(\bar{x}_1) \wedge \cdots \wedge R_m(\bar{x}_m)$ where R is a fixed-point variable. In the case of the above Datalog formula computing the transitive closure of $<$, for example, the rules are $T(x, z) \Leftarrow (x < z)$ and $T(x, z) \Leftarrow T(x, y) \wedge T(y, z)$. When it comes to Datalog-reductions between CSPs, we only need the following intuitive understanding of this concept. A *Datalog-interpretation* is a mapping \mathcal{I} from finite σ -structures to finite τ -structures such that for every finite σ -structure \mathfrak{J} the τ -structure $\mathcal{I}(\mathfrak{J})$ can be defined from \mathfrak{J} or a fixed finite power of it using Datalog formulas [1, Definition 1]. For a τ -structure \mathfrak{A} and σ -structure \mathfrak{A}' , we say that $\text{CSP}(\mathfrak{A}')$ *Datalog-reduces* to $\text{CSP}(\mathfrak{A})$ if there exists a Datalog interpretation \mathcal{I} so that $\mathfrak{J} \rightarrow \mathfrak{A}'$ if and only if $\mathcal{I}(\mathfrak{J}) \rightarrow \mathfrak{A}$. Datalog reductions compose because Datalog interpretations do, i.e., if $\text{CSP}(\mathfrak{A}'')$ Datalog-reduces to $\text{CSP}(\mathfrak{A}')$ and $\text{CSP}(\mathfrak{A}')$ Datalog-reduces to $\text{CSP}(\mathfrak{A})$, then $\text{CSP}(\mathfrak{A}'')$ Datalog-reduces to $\text{CSP}(\mathfrak{A})$.

4 Simplifying the Bodirsky-Pinsker conjecture

In the present section, we give a proof sketch of Theorem 6 and Theorem 8.

4.1 Wreath products

Given groups G and H acting on sets A and B , respectively, their *wreath product* $G \wr H$ is given by $G^B \times H$ acting on $A \times B$ via $((g_b)_{b \in B}, h)(a, b) := (g_{h(b)}(a), h(b))$. The wreath product $G \wr H$ again forms a group, with the neutral element $((e_G)_{b \in B}, e_H)$ and the group operation $((g_b)_{b \in B}, h) \cdot ((g'_b)_{b \in B}, h') := ((g_b \cdot g'_{h^{-1}(b)})_{b \in B}, h \cdot h')$. Note that the permutation group on $A \times B$ induced by the component-wise action of the Cartesian product $G \times H$ is contained by that induced by the action of $G \wr H$ thereon: the permutations induced stemming from the former are represented in $G \wr H$ by all elements of the form $((g)_{b \in B}, h)$.

Below we give an important construction on structures (here also called wreath product) which corresponds to group-theoretic wreath products in terms of their automorphism groups. Let \mathfrak{A} and \mathfrak{B} be structures over disjoint relational signatures τ and σ , and let E be a fresh binary symbol. The *wreath product* $\mathfrak{A} \wr \mathfrak{B}$ is the structure over the signature $\tau \cup \{E\} \cup \sigma$ with domain $A \times B$ whose relations are defined as follows. First, set $E^{\mathfrak{A} \wr \mathfrak{B}} := \{((x_1, y_1), (x_2, y_2)) \in (A \times B)^2 \mid y_1 = y_2\}$. Then, for $R \in \tau$ and $S \in \sigma$ of arities k and m :

$$\begin{aligned} R^{\mathfrak{A} \wr \mathfrak{B}} &:= \{((x_1, y_1), \dots, (x_k, y_k)) \in (A \times B)^k \mid (x_1, \dots, x_k) \in R^{\mathfrak{A}} \text{ and } y_1 = \dots = y_k\}, \\ S^{\mathfrak{A} \wr \mathfrak{B}} &:= \{((x_1, y_1), \dots, (x_m, y_m)) \in (A \times B)^m \mid (y_1, \dots, y_m) \in S^{\mathfrak{B}}\}. \end{aligned}$$

► **Proposition 12** (Basic properties of wreath products). *Let \mathfrak{A} and \mathfrak{B} be structures over disjoint finite relational signatures τ and σ . Then:*

1. $\text{Aut}(\mathfrak{A} \wr \mathfrak{B}) = \text{Aut}(\mathfrak{A}) \wr \text{Aut}(\mathfrak{B})$.
2. *If \mathfrak{A} has no algebraicity, then $\mathfrak{A} \wr \mathfrak{B}$ has no algebraicity.*
3. *If \mathfrak{A} and \mathfrak{B} are ω -categorical, then $\mathfrak{A} \wr \mathfrak{B}$ is ω -categorical.*
4. *If \mathfrak{A} and \mathfrak{B} are homogeneous, then $\mathfrak{A} \wr \mathfrak{B}$ is homogeneous.*
5. *If \mathfrak{A} and \mathfrak{B} are finitely bounded, then $\mathfrak{A} \wr \mathfrak{B}$ is finitely bounded.*
6. *If \mathfrak{A} and \mathfrak{B} are homogeneous Ramsey, then $\mathfrak{A} \wr \mathfrak{B}$ is homogeneous Ramsey.*

Proofs of items 4, 5, and 6 can be found in [33, Lemma 2.2.49] (see also [34, Lemma 2.33]). The statement in item 6 can also be found in [65, Theorem 5.13].

4.2 Removing algebraicity and adding injectivity

In the context of Conjecture 4, the concepts of CSP-injectivity and having no algebraicity are closely related. On the one hand, if \mathfrak{B} is a homogeneous structure with no algebraicity and its relations only contain tuples with pairwise distinct entries, then it is CSP-injective [14, Lemma 4.3.6]. On the other hand, we can prove the following.

► **Proposition 13.** *CSP-injective ω -categorical model-complete cores do not have algebraicity.*

We shall now see that CSP-injectivity and Pol-injectivity cannot hold simultaneously in ω -categorical structures; yet, we shall observe immediately thereafter that CSP-injectivity can be traded against Pol-injectivity using the relation I_4 from Example 5.

► **Proposition 14.**

1. *CSP-injective ω -categorical structures are not Pol-injective.*
2. *The expansion \mathfrak{A}_{I_4} of any CSP-injective structure \mathfrak{A} by the relation I_4 is Pol-injective, and $\text{CSP}(\mathfrak{A})$ and $\text{CSP}(\mathfrak{A}_{I_4})$ are Datalog-interreducible.*

Proof sketch of Theorem 6 and Theorem 8. Let τ and σ be the signatures of \mathfrak{A} and \mathfrak{B} , respectively. We first define the auxiliary structure $\mathfrak{B}' := (\mathbb{Q}; <) \wr \mathfrak{B}$. As shown in Example 3 and Example 5, the structure $(\mathbb{Q}; <)$ is finitely bounded, homogeneous, and has no algebraicity. It then follows immediately from Proposition 12(2) that \mathfrak{B}' has no algebraicity. The expansion \mathfrak{B}'_{\neq} of \mathfrak{B}' by the binary inequality has no algebraicity because taking expansions by first-order definable relations does not change the automorphism group; the same is true for the structure $\mathfrak{B}'_{\neq, I_4}$. Note that removing all symbols in $\sigma \setminus \tau$ as well as the symbol $<$ from $\mathfrak{B}'_{\neq} = ((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ yields the structure $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$. We claim that $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ is CSP-injective. To this end, let \mathfrak{J} be an arbitrary finite structure for which there exists a homomorphism $h: \mathfrak{J} \rightarrow (\mathbb{Q} \wr \mathfrak{A})_{\neq}$. We define \mathfrak{J}_π as the $(\tau \cup \{E\})$ -structure with domain J such that, for every $R \in \tau \cup \{E\}$, we have $\bar{t} \in R^{\mathfrak{J}_\pi}$ if and only if $h(\bar{t}) \in R^{(\mathbb{Q} \wr \mathfrak{A})_{\neq}}$. It is easy to see that \mathfrak{J}_π embeds into the $(\tau \cup \{E\})$ -reduct of \mathfrak{B}' (note that if multiple elements in J are mapped to (q, a) for some $q \in \mathbb{Q}$ and $a \in A$, we can always map them to different elements in $\mathbb{Q} \times \{a\}$). Hence, \mathfrak{J}_π expanded by the binary inequality predicate embeds into $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ and admits an injective homomorphism from \mathfrak{J} . We conclude that $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ is CSP-injective.

Next, we describe two Datalog-reductions \mathcal{I}^* and \mathcal{I} : from $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq})$ to $\text{CSP}(\mathfrak{A})$ and back, starting with the former. Let \mathfrak{J} be an instance of $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq})$. Denote by \sim the equivalence closure of $E^{\mathfrak{J}}$, and let q_\sim be the induced factor map. We define $\mathcal{I}^*(\mathfrak{J})$ as the τ -structure with domain J and whose relations are the preimages of all τ -relations in \mathfrak{J}/\sim under q_\sim , except that we additionally add, for every $x \in J$ such that $(x, x) \in \neq^{\mathfrak{J}}$, the tuple (x, \dots, x) to every τ -relation. It is easy to see that $\mathcal{I}^*(\mathfrak{J})$ is definable in Datalog; we verify

that \mathcal{I}^* is a reduction from $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq})$ to $\text{CSP}(\mathfrak{A})$. Clearly, if \mathfrak{J} maps homomorphically to $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$, then $\mathcal{I}^*(\mathfrak{J})$ maps homomorphically to \mathfrak{A} . Suppose that there exists a homomorphism $h: \mathcal{I}^*(\mathfrak{J}) \rightarrow \mathfrak{A}$. Since by non-triviality \mathfrak{A} does not have a constant endomorphism, there is no $x \in J$ such that $(x, x) \in \neq^{\mathfrak{J}}$. Denote by \mathfrak{J}' the structure obtained from \mathfrak{J} by removing all tuples from $\neq^{\mathfrak{J}}$. Now, choosing representatives j_1, \dots, j_k for all classes of \mathfrak{J}/\sim , the function sending the entirety of $[j_i]_{\sim}$ to $(0, h(j_i))$ for all $i \leq k$ is a homomorphism from \mathfrak{J}' to $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$. By the CSP-injectivity of $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$, it can be replaced by an injective homomorphism, which is then also a homomorphism from \mathfrak{J} to $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$. We continue with the reduction from $\text{CSP}(\mathfrak{A})$ to $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq})$, which is trivial. Let \mathfrak{J} be an instance of $\text{CSP}(\mathfrak{A})$. We define $\mathcal{I}(\mathfrak{J})$ as the $(\tau \cup \{E, \neq\})$ -expansion of \mathfrak{J} by empty relations. By CSP-injectivity, there exists a homomorphism $h: \mathcal{I}(\mathfrak{J}) \rightarrow (\mathbb{Q} \wr \mathfrak{A})_{\neq}$ if and only if there exists an injective such homomorphism i . This is the case if and only if $E^{(\mathbb{Q} \wr \mathfrak{A})_{\neq}}$ restricted to $i(J)$ describes the kernel of a homomorphism from \mathfrak{J} to \mathfrak{A} . Therefore, $\mathfrak{J} \rightarrow \mathfrak{A}$ if and only if $\mathcal{I}(\mathfrak{J}) \rightarrow (\mathbb{Q} \wr \mathfrak{A})_{\neq}$.

Now consider $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$. By Proposition 14(2), $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$ is Pol-injective and $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4})$ and $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq})$ are Datalog-interreducible. That $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4})$ and $\text{CSP}(\mathfrak{A})$ are Datalog-interreducible then follows from the fact that Datalog-reductions can be composed.

Finally, we sketch how to verify the properties in items 1, 2, 3, and 4. We only explicitly cover the case where I_4 is present (Theorem 8), since the arguments without I_4 are a subset.

Regarding item 1, the first part (ω -categoricity) follows directly from Proposition 12(3) applied to \mathfrak{B}' and since we take an expansion by relations which are preserved by all bijections, and in particular all automorphisms. The second part of item 1 (model-complete cores) is immediate once one observes that the endomorphism monoid of a wreath product admits a similar description as its automorphism group, cf. Proposition 12(1).

Regarding item 2 (homogeneity and the Ramsey property), one first observes that, since $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq, I_4}$ is an expansion of \mathfrak{B}' by two relations that are preserved by all embeddings, we can ignore these relations and only prove the statement for \mathfrak{B}' . Then homogeneity follows directly from Proposition 12(4) because $(\mathbb{Q}; <)$ is homogeneous. In the second part, the backward direction follows directly from Proposition 12(6) because $(\mathbb{Q}; <)$ is homogeneous Ramsey [53]. The forward direction can be proved by hand, by instantiating the Ramsey property for the finite substructures of \mathfrak{B}' where E interprets as the diagonal relation.

Regarding item 3 (finite boundedness), this is a consequence of Proposition 12(5) (which immediately yields finite boundedness of \mathfrak{B}') and the fact that we take an expansion by relations definable by a Boolean combination of equality atoms. More precisely, recall that up to isomorphism, the finite substructures of a finitely bounded structure are precisely the finite models of a universal sentence. By adding new clauses defining the relations \neq and I_4 in terms of equalities to the universal sentence describing the finite substructures of \mathfrak{B}' , we obtain a universal sentence describing the finite substructures of $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq, I_4}$.

Regarding item 4 (pp-construction of \mathfrak{K}_3), we use a characterization via the pseudo-Siggers identity from [7, Theorems 1.3 and 3.4] and the fact that $(\mathbb{Q}; \neq, I_4)$ is an ω -categorical model-complete core that does not pp-construct \mathfrak{K}_3 [14, Theorems 12.0.1, 12.7.3, 12.9.2, and Corollary 6.4.4]. One direction is simple (and does not use the assumption on the orbit growth), because $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$ pp-constructs \mathfrak{A} ; the pp-construction is given by $\mathfrak{A} = ((\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4} / E^{(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}})^{\tau}$ (factors and reducts are pp-constructions [9]). Thus, if \mathfrak{A} pp-constructs \mathfrak{K}_3 , then so does $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$. For the other direction, one assumes that \mathfrak{A} does not pp-construct \mathfrak{K}_3 , in which case Theorem 1.3 in [7] yields polymorphisms in its model-complete core satisfying the pseudo-Siggers identity. These are then composed with polymorphisms of $(\mathbb{Q}; \neq, I_4)$ witnessing the pseudo-Siggers identity, yielding the satisfaction of the pseudo-Siggers identity by polymorphisms of the model-complete core of $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$. The proof of

Proposition 12(3) shows that also the structure $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$ has less than double exponential orbit growth. It then follows from Theorem 3.4 in [7] (whose item (ii) is equivalent to the pp-construction of \mathfrak{K}_3 by Theorem 1.8 in [9]) that $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$ does not pp-construct \mathfrak{K}_3 . ◀

► **Example 15.** Consider the two structures $\mathfrak{A}_1 := (\{0, 1\}; \{1\}, \{(x, y, z) \in \{0, 1\}^3 \mid x + y + z = 0 \pmod{2}\})$ and $\mathfrak{A}_2 := (\mathbb{Q}; <, \{(x, y, z) \in \mathbb{Q}^3 \mid x \geq y \text{ or } x \geq z\})$. The former is a finitely bounded homogeneous model-complete core *with* algebraicity (due to it being finite) and the latter is a finitely bounded homogeneous model-complete core *without* algebraicity. Note that \mathfrak{A}_1 is preserved by the symmetric operation $(x_1, \dots, x_n) \mapsto x_1 + \dots + x_n \pmod{2}$ for every odd $n \geq 2$ and \mathfrak{A}_2 is preserved by the symmetric operation $(x_1, \dots, x_n) \mapsto \min\{x_1, \dots, x_n\}$ for every $n \geq 2$. Thus, in both cases the polymorphism clone satisfies the cyclic identity of arity n for some $n \geq 2$. Since pp-constructions preserve the satisfaction of height-1 identities in polymorphism clones [9] and no cyclic operation f of arity $n \geq 2$ preserves \neq over an infinite set, \mathfrak{A}_i does not pp-construct $(\mathbb{Q} \wr \mathfrak{A}_i)_{\neq}$ for both $i \in [2]$. Hence, our Datalog-reduction in Theorem 6 is not subsumed by gadget reductions (pp-constructions).

5 ω -categorical cheeses for PCSPs

In the present section, we give a proof sketch of Theorem 10.

5.1 Full powers

Our basic tool for creating ω -categorical cheeses for PCSPs that are not finitely tractable are *full powers* (see, e.g., Bodirsky [14, Section 3.5]). Roughly speaking, a full power $\mathfrak{A}^{[d]}$ is a higher-dimensional representation of \mathfrak{A} that can be pp-constructed from it and vice versa; its central feature is that its polymorphisms are the polymorphisms of \mathfrak{A} acting on d -tuples, thus allowing to factor it by the orbit-equivalence of $\text{Aut}(\mathfrak{A})$ on d -tuples (rather than on elements). In other words, the orbits of d -tuples of \mathfrak{A} are represented by orbits of 1-tuples of $\mathfrak{A}^{[d]}$, which is crucial in applications of the sandwich method. Let \mathfrak{A} be a relational structure with a signature τ and let $d \in \mathbb{N}$ be arbitrary. The d -th *full power* of \mathfrak{A} , denoted $\mathfrak{A}^{[d]}$, is the structure with domain A^d and the following relations for every $k \in [d]$:

- for every $R \in \tau$ of arity k and every injection $\iota: [k] \rightarrow [d]$, the unary relation

$$R_\iota^{\mathfrak{A}^{[d]}} := \{(a_1, \dots, a_d) \in A^d \mid (a_{\iota(1)}, \dots, a_{\iota(k)}) \in R^{\mathfrak{A}}\};$$

- for every $R \in \tau$ of arity k and every function $\iota: [k] \rightarrow [d]$, the k -ary relation

$$\widehat{R}_\iota^{\mathfrak{A}^{[d]}} := \{((a_1^1, \dots, a_d^1), \dots, (a_1^k, \dots, a_d^k)) \in (A^d)^k \mid (a_{\iota(1)}^1, \dots, a_{\iota(k)}^k) \in R^{\mathfrak{A}}\};$$

- for all functions $\iota, \iota': [k] \rightarrow [d]$, the binary *compatibility relation*

$$S_{\iota, \iota'}^{\mathfrak{A}^{[d]}} := \{((a_1^1, \dots, a_d^1), (a_1^2, \dots, a_d^2)) \in (A^d)^2 \mid \forall i \in [k]: a_{\iota(i)}^1 = a_{\iota'(i)}^2\}.$$

Below we give several important properties of full powers.

► **Proposition 16** (Basic properties of full powers). *Let \mathfrak{A} be a relational structure over a finite signature τ and let its relations be of arity $\leq d$. Then:*

1. $\text{Pol}(\mathfrak{A}^{[d]})$ consists of the component-wise actions of polymorphisms of \mathfrak{A} on d -tuples.
2. $\mathfrak{A}^{[d]}$ is (factor-free) pp-interpretable from \mathfrak{A} and vice versa.
3. If \mathfrak{A} is ω -categorical, then $\mathfrak{A}^{[d]}$ is ω -categorical.
4. If \mathfrak{A} is homogeneous, then $\mathfrak{A}^{[d]}$ is homogeneous.

5. If \mathfrak{A} is a model-complete core, then $\mathfrak{A}^{[d]}$ is a model-complete core.
6. If \mathfrak{A} is finitely bounded, then $\mathfrak{A}^{[d]}$ is finitely bounded.
7. If \mathfrak{A} is homogeneous Ramsey, then $\mathfrak{A}^{[d]}$ is homogeneous Ramsey.
8. If \mathfrak{A} is a reduct of \mathfrak{B} , then $\mathfrak{A}^{[d]}$ is a reduct of $\mathfrak{B}^{[d]}$.

5.2 Our sandwich recipe

To construct a finite PCSP sandwich for a structure \mathfrak{A} , we must only select a finite substructure \mathfrak{S}_1 and a finite factor \mathfrak{S}_2 of \mathfrak{A} . The issue is that, in many cases, $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ will be finitely tractable, possibly because already $\text{CSP}(\mathfrak{S}_1)$ or $\text{CSP}(\mathfrak{S}_2)$ is tractable. Proposition 17 describes a general condition for structures in the scope of Conjecture 4 under which this does not happen, i.e., where $\text{PCSP}(\mathfrak{S}_1, \mathfrak{S}_2)$ is provably not finitely tractable. Intuitively, in order to obtain a non-finitely tractable sandwich for a structure in the scope of Conjecture 4, it is enough to take a sufficiently large full power thereof, a finite substructure \mathfrak{S}_1 whose polymorphisms preserve the inequality, and a finite factor \mathfrak{S}_2 by an arbitrary subgroup of automorphisms preserving a linear order and acting with finitely many orbits.

► **Proposition 17.** *Let \mathfrak{A} be a reduct of a linearly ordered finitely bounded homogeneous structure \mathfrak{B} such that one of the relations of \mathfrak{A} interprets as \neq . Let $d \in \mathbb{N}$ be such that the bounds of \mathfrak{B} have size $\leq d$ and its relations are of arity $\leq d - 1$. Then, for every finite substructure \mathfrak{S} of \mathfrak{A} with $|S| \geq 3$:*

1. $\mathfrak{A}^{[d]}$ is a cheese for $\text{PCSP}(\mathfrak{S}^{[d]}, \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]})$;
2. Every finite cheese for $\text{PCSP}(\mathfrak{S}^{[d]}, \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]})$ pp-constructs \mathfrak{K}_3 .

Proof sketch. Clearly, for every finite substructure \mathfrak{S} of \mathfrak{A} , we have $\mathfrak{S}^{[d]} \rightarrow \mathfrak{A}^{[d]} \rightarrow \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]}$. Now suppose that there is a finite cheese \mathfrak{D} of $\text{PCSP}(\mathfrak{S}^{[d]}, \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]})$ that does not pp-construct \mathfrak{K}_3 . Then, by Theorem 11, there exists a cyclic operation f in $\text{Pol}(\mathfrak{D})$ of some arity n . Take any homomorphisms $g: \mathfrak{S}^{[d]} \rightarrow \mathfrak{D}$ and $h: \mathfrak{D} \rightarrow \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]}$. Then, by setting $f'(x_1, \dots, x_n) := h \circ f(g(x_1), \dots, g(x_n))$, we get a cyclic $f' \in \text{Pol}(\mathfrak{S}^{[d]}, \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]})$ of arity n . One can show that every polymorphism from $\mathfrak{S}^{[d]}$ to $\mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]}$ can be lifted to a polymorphism from \mathfrak{S} to \mathfrak{A} and that any height-1 identity that is satisfied in $\text{Pol}(\mathfrak{S}^{[d]}, \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]})$ is also satisfied in $\text{Pol}(\mathfrak{S}, \mathfrak{A})$ modulo $\text{Aut}(\mathfrak{B})$. Hence, there is some n -ary $g \in \text{Pol}(\mathfrak{S}, \mathfrak{A})$ that is pseudo-cyclic modulo $\text{Aut}(\mathfrak{B})$. One easily sees that any operation that is pseudo-cyclic modulo some order-preserving automorphism group must be cyclic. This means that we have successfully lifted the cyclic identity from $\text{Pol}(\mathfrak{S}^{[d]}, \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]})$ to $\text{Pol}(\mathfrak{S}, \mathfrak{A})$. But any cyclic operation on a set of size at least 3 does not preserve \neq , a contradiction. ◀

The prerequisites of Proposition 17 are fairly general but clearly not sufficient to cover the entire scope of Conjecture 4 (cf. Section 6.2); this is where Theorem 6 becomes handy. Note that removing algebraicity using Theorem 6 allows us to expand our structure by the inequality relation without fundamentally changing the complexity of its CSP. It also allows us to identify a subgroup of automorphisms preserving a linear order and acting with finitely many orbits, via the following folklore result (see e.g. [14, Section 2.3.6]).

► **Proposition 18.** *Let $\mathfrak{A}_1, \mathfrak{A}_2$ be countable homogeneous structures without algebraicity over disjoint finite relational signatures τ_1, τ_2 . Then there exists an up to isomorphism unique countable homogeneous $(\tau_1 \cup \tau_2)$ -structure $\mathfrak{A}_1 \otimes \mathfrak{A}_2$, called the generic superposition of $\mathfrak{A}_1, \mathfrak{A}_2$, such that any finite $(\tau_1 \cup \tau_2)$ -structure embeds into $\mathfrak{A}_1 \otimes \mathfrak{A}_2$ if and only if its τ_i -reduct embeds into \mathfrak{A}_i for both $i \in \{1, 2\}$. Moreover, if $\mathfrak{A}_1, \mathfrak{A}_2$ are both finitely bounded, then so is $\mathfrak{A}_1 \otimes \mathfrak{A}_2$.*

Proof sketch (Theorem 10). Let \mathfrak{A} be a non-trivial reduct of a finitely bounded homogeneous structure \mathfrak{B} . In the first step, we move to the structures without algebraicity $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ and $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ from Theorem 6; recall that $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ pp-constructs \mathfrak{K}_3 if and only if \mathfrak{A} does and that $\text{CSP}(\mathfrak{A})$ and $\text{CSP}((\mathbb{Q} \wr \mathfrak{A})_{\neq})$ are Datalog-interreducible. Next, we take the generic superposition of $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ with $(\mathbb{Q}; <)$, which exists by Proposition 18. We set $\widehat{\mathfrak{B}} := ((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq} \otimes (\mathbb{Q}; <)$ and define $\widehat{\mathfrak{A}}$ as the reduct of $\widehat{\mathfrak{B}}$ to the signature of $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$. Then $\widehat{\mathfrak{A}}$ and $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ are isomorphic, and $\widehat{\mathfrak{B}}$ is finitely bounded homogeneous due to Proposition 18. But now $\widehat{\mathfrak{A}}$ and $\widehat{\mathfrak{B}}$ satisfy the prerequisites of Proposition 17. Let \mathcal{I} be the composition of the (trivial) Datalog reduction \mathcal{I} from \mathfrak{A} to $(\mathbb{Q} \wr \mathfrak{A})_{\neq} \cong \widehat{\mathfrak{A}}$ from Theorem 6 and the (factor-free) pp-interpretation of $\widehat{\mathfrak{A}}^{[d]}$ from $\widehat{\mathfrak{A}}$ and let \mathcal{J}' be the composition of the pp-interpretation and Datalog reduction running in the opposite direction.

We select a finite substructure \mathfrak{S} of \mathfrak{A} with $|S| \geq 3$ large enough so that every finite structure with a homomorphism to a substructure of \mathfrak{A} of size $\leq n$ also has a homomorphism to \mathfrak{S} ; let $\widehat{\mathfrak{S}}$ be the substructure of $\widehat{\mathfrak{A}}$ induced by $\{0\} \times S$. We claim that $\mathfrak{A}' := \widehat{\mathfrak{A}}^{[d]}$, $\mathfrak{B}' := \widehat{\mathfrak{B}}^{[d]}$, $\mathfrak{S}_1 := \widehat{\mathfrak{S}}^{[d]}$, and $\mathfrak{S}_2 := \widehat{\mathfrak{A}}^{[d]} \widehat{\mathfrak{B}}^{[d]}$ for d as in Proposition 17 witness items 1–7 of the theorem. Items 1 to 6 follow from Theorem 6, Proposition 16, and Proposition 17. For item 7, note that if \mathfrak{X} maps homomorphically to a substructure of \mathfrak{A} of size $\leq n$, then $\mathfrak{X} \rightarrow \mathfrak{S}$, hence $\mathcal{I}(\mathfrak{X}) \rightarrow \widehat{\mathfrak{S}}$ and hence $\mathcal{J}'(\mathfrak{X}) = \mathcal{I}(\mathfrak{X})^{[d]} \rightarrow \widehat{\mathfrak{S}}^{[d]} = \mathfrak{S}_1$. Conversely, if $\mathfrak{X} \twoheadrightarrow \mathfrak{A}$, then $\mathcal{I}(\mathfrak{X}) \twoheadrightarrow \widehat{\mathfrak{A}}$ by Theorem 6. One can show that if $\mathcal{I}(\mathfrak{X})^{[d]} \rightarrow \mathfrak{S}_2$, then also $\mathcal{I}(\mathfrak{X}) \twoheadrightarrow \widehat{\mathfrak{A}}$, hence $\mathcal{J}'(\mathfrak{X}) \twoheadrightarrow \mathfrak{S}_2$. ◀

6 Conclusion and outlook

6.1 Topology seems irrelevant

In the context of infinite-domain CSPs, the influence of topological properties of the polymorphism clones of ω -categorical structures has received much attention, with various results claiming their relevance [26] or irrelevance [11]: while generally, whether or not an ω -categorical structure \mathfrak{A} pp-constructs \mathfrak{K}_3 (or pp-interprets another ω -categorical structure \mathfrak{A}') does depend on such topological properties [9, 29, 47, 31], in many situations it does not (see, e.g., [42]). By moving from any ω -categorical model-complete core template to one without algebraicity, Theorem 6 allows us to restrict Conjecture 4 to a class of topologically well-behaved structures. Namely, for those recent research suggests that often the algebraic structure determines the topological one; at least this is a fact for the endomorphism monoid, where the relevant topology (*pointwise convergence*) can then be defined from the algebraic structure of the monoid alone [63]. For ω -categorical structures with algebraicity, on the other hand, various examples [43, 17, 63] show that this need not be the case. This also implies that there is no hope of lifting isomorphisms between endomorphism monoids or polymorphism clones of ω -categorical structures to their counterparts constructed in Theorem 6.

6.2 Necessity of cheese preprocessing

In our proof of Theorem 10, we apply Proposition 17 to a structure which we previously “preprocessed” using Theorem 6 (removing algebraicity) and generic superpositions. It is natural to ask how strong Proposition 17 is on its own: can a result similar to Theorem 10 be proved directly using this proposition, perhaps under a more general structural assumption, e.g. that \mathfrak{A} is a model-complete core? Here the answer seems to be negative. Consider the model-complete core structure $\mathfrak{A} = (\mathbb{Q}; <, \{(x, y, z) \in \mathbb{Q}^3 \mid x \geq y \text{ or } x \geq z\})$ within the scope of Conjecture 4 (\mathfrak{A}_2 in Example 15); its CSP is polynomial-time tractable [22]. Let \mathfrak{S} be any finite structure for which there exists a homomorphism $h: \mathfrak{S} \rightarrow \mathfrak{A}^{[d]}$, and let \mathfrak{B} be an

arbitrary ω -categorical expansion of \mathfrak{A} . Let A_S be the set of the elements in A that appear as an entry of some element of $h(S)$, and let \mathfrak{A}_S be the substructure of \mathfrak{A} on A_S . Since $\mathfrak{A}^{[d]} \rightarrow \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]}$, the structure $\mathfrak{A}_S^{[d]}$ is a finite cheese for $\text{PCSP}(\mathfrak{S}, \mathfrak{A}_{/\text{Aut}(\mathfrak{B})}^{[d]})$.

We claim that $\text{CSP}(\mathfrak{A}_S^{[d]})$ is tractable, and hence this PCSP is finitely tractable. To see this, observe that the binary minimum operation $\min(x, y)$ is a polymorphism of \mathfrak{A} . This operation is *conservative*: $\min(x, y) \subseteq \{x, y\}$ holds for all x, y . Hence, its restriction to A_S^2 induces a polymorphism of \mathfrak{A}_S , which itself induces a polymorphism of $\mathfrak{A}_S^{[d]}$ through its component-wise action. Since this operation is cyclic, $\mathfrak{A}_S^{[d]}$ does not pp-construct \mathfrak{K}_3 [9, Theorem 1.4], and hence we are done by Theorem 1. This issue cannot be fixed simply by taking an expansion of \mathfrak{A} by the inequality relation \neq because $\text{CSP}(\mathfrak{A}; \neq)$ is NP-complete [22].

6.3 Algorithmic properties of our sandwiches

An interesting open question is whether the PCSP templates generated by Theorem 10 can be solved by some universal algorithm. Many known PCSP problems admitting infinite tractable cheeses can be solved by numeric relaxation algorithms such as *BLP* (the basic linear programming relaxation), *AIP* (the basic affine integer relaxation) or *BLP+AIP* (a combination of both) (see, e.g., [6, 35]). The proof of Proposition 17 shows that none of the PCSP templates produced by Theorem 10 admits a cyclic polymorphism. Since solvability of PCSPs by BLP is characterized by admitting totally symmetric polymorphisms of all arities [6, Theorem 7.9], which would also be cyclic, BLP does not solve any of those PCSPs. Regarding the combination BLP+AIP, we remark that replacing $(\mathbb{Q} \wr \mathfrak{A})_{\neq}$ and $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq}$ in the proof of Theorem 10 by $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$ and $((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq, I_4}$ ensures that there is no cyclic and no 2-block symmetric operation in $\text{Pol}(\mathfrak{S}, (\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4})$. In fact, we get the following much stronger statement, which follows directly from the fact that we take an expansion by I_4 .

► **Proposition 19** (cf. proof of Lemma 7.5.1 in [14]). *Let Σ be any finite height-1 condition that cannot be satisfied by any set of essentially injective operations from an at least 2-element set S_1 to a countably infinite set S_2 . Then, for every at least 2-element substructure \mathfrak{S} of $(\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$, we have that $\text{Pol}(\mathfrak{S}, (\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4})$ does not satisfy Σ .* ◀

Similarly to cyclic identities, also 2-block symmetric identities can be lifted from $\text{Pol}(\mathfrak{S}^{[d]}, \widehat{\mathfrak{A}}_{/G}^{[d]})$ to $\text{Pol}(\mathfrak{S}, \widehat{\mathfrak{A}})$, where $\widehat{\mathfrak{A}} := (\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$ and $G := \text{Aut}(((\mathbb{Q}; <) \wr \mathfrak{B})_{\neq, I_4} \otimes (\mathbb{Q}; <))$. Thus, since solvability of a finite-domain PCSP by BLP+AIP is characterized by the existence of 2-block symmetric polymorphisms [35, Theorem 4], we get that $\text{PCSP}(\mathfrak{S}^{[d]}, \widehat{\mathfrak{A}}_{/G}^{[d]})$ is not solvable by BLP+AIP. For details, we refer the reader to [58, Theorem 2], where this idea was used in concrete cases of non-finitely tractable PCSPs. It follows that we can exclude BLP and BLP+AIP as potential algorithms solving all tractable PCSPs produced by Theorem 10. As demonstrated by Mottet [58, Lemma 33], the above idea can in fact be extended to any universal algorithm such that solvability of finite-domain PCSPs by this algorithm is captured by a set of height-1 identities specified by non-trivial permutations of variables. It is folklore that most universal algorithms for finite-domain PCSPs are either captured by such height-1 conditions or their algebraic description is rather complicated [38].

Open question: Is there a universal algorithm that would solve all tractable PCSPs stemming from Theorem 10, even when $\widehat{\mathfrak{A}} = (\mathbb{Q} \wr \mathfrak{A})_{\neq}$ is replaced by $\widehat{\mathfrak{A}} = (\mathbb{Q} \wr \mathfrak{A})_{\neq, I_4}$?

The lifting method we and Mottet [58] use provably cannot be extended to all height-1 conditions covered by Proposition 19. To see this, note that the Olšák identities $f(y, x, x, x, y, y) \approx f(x, y, x, x, y, y) \approx f(x, x, y, y, y, x)$ cannot be satisfied by essentially injective operations but every finite-domain PCSP whose template does not have an Olšák

polymorphism must be NP-hard [6, Corollary 6.3]. This means that the Olšák identities cannot possibly be prevented in our PCSPs without losing solvability in polynomial time (unless $P=NP$).

References

- 1 Albert Atserias, Andrei A. Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. *Theor. Comput. Sci.*, 410(18):1666–1683, 2009. doi:10.1016/j.tcs.2008.12.049.
- 2 Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2 + \varepsilon)$ -Sat Is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017. A conference version appeared in the proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science - FOCS'14. doi:10.1137/15M1006507.
- 3 Franz Baader and Jakub Rydval. Using Model Theory to Find Decidable and Tractable Description Logics with Concrete Domains. *J. Autom. Reason.*, 66:357–407, 2022. doi:10.1007/s10817-022-09626-2.
- 4 Libor Barto. Promises Make Finite (Constraint Satisfaction) Problems Infinitary. In *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science - LICS'19*, pages 1–8. IEEE, 2019. doi:10.1109/lics.2019.8785671.
- 5 Libor Barto, Bertalan Bodor, Marcin Kozik, Antoine Mottet, and Michael Pinsker. Symmetries of graphs and structures that fail to interpret a finite thing. In *Proceedings of the 38th Annual ACM/IEEE Symposium on Logic in Computer Science - LICS'23*, pages 1–13. IEEE, 2023. doi:10.1109/LICS56636.2023.10175732.
- 6 Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM.*, 68(4):28:1–28:66, 2021. doi:10.1145/3457606.
- 7 Libor Barto, Michael Kompatscher, Miroslav Olšák, Van Pham Trung, and Michael Pinsker. Equations in oligomorphic clones and the constraint satisfaction problem for ω -categorical structures. *J. Math. Log.*, 19(02):Article 1950010, 2019. doi:10.1142/S0219061319500107.
- 8 Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Log. Methods Comput. Sci.*, 8(1):7:1–7:27, 2012. doi:10.2168/lmcs-8(1:7)2012.
- 9 Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. *Isr. J. Math.*, 223(1):363–398, 2018. doi:10.1007/s11856-017-1621-9.
- 10 Libor Barto and Michael Pinsker. The algebraic dichotomy conjecture for infinite domain constraint satisfaction problems. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science - LICS'16*, pages 615–622. ACM, 2016. doi:10.1145/2933575.2934544.
- 11 Libor Barto and Michael Pinsker. Topology is irrelevant (in a dichotomy conjecture for infinite domain constraint satisfaction problems). *SIAM J. Comput.*, 49(2):365–393, 2020. doi:10.1137/18M1216213.
- 12 Zeno Bitter and Antoine Mottet. Generalized completion problems with forbidden tournaments. In *Proceedings of the 49th International Symposium on Mathematical Foundations of Computer Science, MFCS'24*, volume 306 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.MFCS.2024.28.
- 13 Manuel Bodirsky. Cores of countably categorical structures. *Log. Methods Comput. Sci.*, 3(1):2:1–2:16, 2007. doi:10.2168/LMCS-3(1:2)2007.
- 14 Manuel Bodirsky. *Complexity of infinite-domain constraint satisfaction*, volume 52. Cambridge University Press, 2021. Postprint available online. doi:10.1017/9781107337534.
- 15 Manuel Bodirsky, Hubie Chen, and Michael Pinsker. The reducts of equality up to primitive positive interdefinability. *J. Symb. Log.*, 75(4):1249–1292, 2010. doi:10.2178/js1/1286198146.

- 16 Manuel Bodirsky and Víctor Dalmau. Datalog and constraint satisfaction with infinite templates. *J. Comput. Syst. Sci.*, 79(1):79–100, 2013. A conference version appeared in the proceedings of the 34th Symposium on Theoretical Aspects of Computer Science - STACS’06. doi:10.1016/j.jcss.2012.05.012.
- 17 Manuel Bodirsky, David M. Evans, Michael Kompatscher, and Michael Pinsker. A counterexample to the reconstruction of ω -categorical structures from their endomorphism monoid. *Isr. J. Math.*, 224:57–82, 2018. doi:10.1007/s11856-018-1645-9.
- 18 Manuel Bodirsky and Johannes Greiner. The complexity of combinations of qualitative constraint satisfaction problems. *Log. Methods Comput. Sci.*, 16(1):21:1–21:19, 2020. doi:10.23638/LMCS-16(1:21)2020.
- 19 Manuel Bodirsky and Johannes Greiner. Tractable combinations of theories via sampling. In *Proceedings of the 17th European Conference on Logics in Artificial Intelligence - JELIA’23*, pages 133–146. Springer, 2021. doi:10.1007/978-3-030-75775-5_10.
- 20 Manuel Bodirsky and Martin Grohe. Non-dichotomies in constraint satisfaction complexity. In *Proceedings of the 35th International Colloquium on Automata, Languages, and Programming - ICALP’08*, pages 184–196. Springer, 2008. doi:10.1007/978-3-540-70583-3_16.
- 21 Manuel Bodirsky, Peter Jonsson, and Trung Van Pham. The complexity of phylogeny constraint satisfaction problems. *ACM Transactions on Computational Logic (TOCL)*, 18(3):1–42, 2017. doi:10.1145/3105907.
- 22 Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *J. ACM.*, 57(2):9:1–9:41, 2010. A conference version appeared in the proceedings of the 40th ACM Symposium on Theory of Computing - STOC’08. doi:10.1145/1667053.1667058.
- 23 Manuel Bodirsky, Florent R. Madelaine, and Antoine Mottet. A proof of the algebraic tractability conjecture for monotone monadic SNP. *SIAM J. Comput.*, 50(4):1359–1409, 2021. doi:10.1137/19M128466X.
- 24 Manuel Bodirsky, Barnaby Martin, Michael Pinsker, and András Pongrácz. Constraint satisfaction problems for reducts of homogeneous graphs. *SIAM J. Comput.*, 48(4):1224–1264, 2019. A conference version appeared in the proceedings of the 43rd International Colloquium on Automata, Languages, and Programming, ICALP’16. doi:10.1137/16M1082974.
- 25 Manuel Bodirsky and Antoine Mottet. Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science - LICS’16*, pages 623–632. ACM, 2016. doi:10.1145/2933575.2934515.
- 26 Manuel Bodirsky, Antoine Mottet, Miroslav Olšák, Jakub Opršal, Michael Pinsker, and Ross Willard. Topology is relevant (in a dichotomy conjecture for infinite-domain constraint satisfaction problems). In *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science - LICS’19*, pages 1–12. IEEE, 2019. doi:10.1109/LICS.2019.8785883.
- 27 Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *J. Log. Comput.*, 16(3):359–373, 2006. A conference version appeared in the proceedings of Computer Science Logic - CSL’03. doi:10.1093/logcom/exi083.
- 28 Manuel Bodirsky and Michael Pinsker. Schaefer’s theorem for graphs. *Journal of the ACM*, 62(3):19:1–19:52, 2015. A conference version appeared in the proceedings of the 43rd ACM Symposium on Theory of Computing - STOC’11. doi:10.1145/2764899.
- 29 Manuel Bodirsky and Michael Pinsker. Topological Birkhoff. *Trans. Amer. Math. Soc.*, 367(4):2527–2549, 2015. doi:10.1090/S0002-9947-2014-05975-8.
- 30 Manuel Bodirsky and Michael Pinsker. Canonical functions: a proof via topological dynamics. *Contrib. Discrete Math.*, 16(2):36–45, 2021. doi:10.11575/cdm.v16i2.71724.
- 31 Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. *J. Symb. Log.*, 86(1):148–161, 2021. doi:10.1017/jsl.2019.23.
- 32 Manuel Bodirsky and Jakub Rydval. On the descriptive complexity of temporal constraint satisfaction problems. *J. ACM.*, 70(1):2:1–2:58, 2022. doi:10.1145/3566051.

- 33 Bertalan Bodor. *CSP dichotomy for ω -categorical monadically stable structures*. Phd thesis, Technische Universität Dresden, 2021.
- 34 Bertalan Bodor. Classification of-categorical monadically stable structures. *The Journal of Symbolic Logic*, 89(2):460–495, 2024. doi:10.1017/jsl.2023.66.
- 35 Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. The power of the combined basic linear programming and affine relaxation for promise constraint satisfaction problems. *SIAM J. Comput.*, 49(6):1232–1248, 2020. doi:10.1137/20M1312745.
- 36 Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science - FOCS'17*, pages 319–330. IEEE, 2017. doi:10.1109/FOCS.2017.37.
- 37 Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005. doi:10.1137/S0097539700376676.
- 38 Lorenzo Ciardo and Stanislav Živný. Hierarchies of minion tests for PCSPs through tensors. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms - SODA'23*, pages 568–580. SIAM, 2023. doi:10.1137/1.9781611977554.ch25.
- 39 Victor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science - LICS'24*, pages 1–15, 2024. doi:10.1145/3661814.3662068.
- 40 Guofeng Deng, Ezzeddine El Sai, Trevor Manders, Peter Mayr, Poramate Nakkirt, and Athena Sparks. Sandwiches for promise constraint satisfaction. *Algebra Univers.*, 82(1):15:1–15:8, 2021. doi:10.1007/s00012-020-00702-5.
- 41 Irit Dinur, Oded Regev, and Clifford Smyth. The hardness of 3-uniform hypergraph coloring. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 33–40. IEEE, 2002. doi:10.1109/SFCS.2002.1181880.
- 42 L. Elliott, J. Jonušas, J.D. Mitchell, Y. Péresse, and M. Pinsker. Polish topologies on endomorphism monoids of relational structures. *Advances in Mathematics*, 431:Article 109214, 2023. doi:10.1016/j.aim.2023.109214.
- 43 David M. Evans and Paul R. Hewitt. Counterexamples to a conjecture on relative categoricity. *Ann. Pure Appl. Log.*, 46(2):201–209, 1990. doi:10.1016/0168-0072(90)90034-Y.
- 44 Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1998. doi:10.1137/S0097539794266766.
- 45 Roman Feller and Michael Pinsker. An algebraic proof of the graph orientation problem dichotomy for forbidden tournaments. Preprint arXiv:2405.20263, 2024. doi:10.48550/arXiv.2405.20263.
- 46 Michael R. Garey and David S. Johnson. The complexity of near-optimal graph coloring. *J. ACM.*, 23(1):43–49, 1976. doi:10.1145/321921.321926.
- 47 Mai Gehrke and Michael Pinsker. Uniform Birkhoff. *J. Pure Appl. Algebra*, 222(5):1242–1250, 2018. doi:10.1016/j.jpaa.2017.06.016.
- 48 Pierre Gillibert, Julius Jonusas, Michael Kompatscher, Antoine Mottet, and Michael Pinsker. When symmetries are not enough: A hierarchy of hard constraint satisfaction problems. *SIAM J. Comput.*, 51(2):175–213, 2022. doi:10.1137/20M1383471.
- 49 Pierre Gillibert, Julius Jonušas, Michael Kompatscher, Antoine Mottet, and Michael Pinsker. Hrushovski's encoding and ω -categorical CSP monsters. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming- ICALP'20*, volume 168 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 131:1–131:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.131.
- 50 Jan Hubička and Jaroslav Nešetřil. All those Ramsey classes (Ramsey classes with closures and forbidden homomorphisms). *Adv. Math.*, 356:Article 106791, 2019. doi:10.1016/j.aim.2019.106791.

- 51 Peter Jeavons. On the algebraic structure of combinatorial problems. *Theor. Comput. Sci.*, 200(1-2):185–204, 1998. doi:10.1016/S0304-3975(97)00230-2.
- 52 Peter Jeavons, David Cohen, and Marc Gyssens. Closure properties of constraints. *J. ACM.*, 44(4):527–548, 1997. doi:10.1145/263867.263489.
- 53 Alexander S. Kechris, Vladimir G. Pestov, and Stevo Todorćević. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geom. Funct. Anal.*, 15:106–189, 2005. doi:10.1007/s00039-005-0503-1.
- 54 Michael Kompatscher and Trung Van Pham. A complexity dichotomy for poset constraint satisfaction. In *Proceedings of the 34th Symposium on Theoretical Aspects of Computer Science - STACS'17*, volume 66 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.STACS.2017.47.
- 55 Dugald Macpherson. A survey of homogeneous structures. *Discrete Math.*, 311(15):1599–1634, 2011. doi:10.1016/j.disc.2011.01.024.
- 56 Paolo Marimon and Michael Pinsker. Binary symmetries of tractable non-rigid structures (Minimal operations over permutation groups). To appear at LICS'25; preprint arXiv:2410.22060, 2025.
- 57 Antoine Mottet. Promise and infinite-domain constraint satisfaction. In *Proceedings of the 32nd EACSL Annual Conference on Computer Science Logic - CSL'24*, volume 288 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 41:1–41:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.CSL.2024.41.
- 58 Antoine Mottet. Algebraic and algorithmic synergies between promise and infinite-domain CSPs. to appear at LICS'25; preprint arXiv:2501.13740, 2025. doi:10.48550/arXiv.2501.13740.
- 59 Antoine Mottet, Tomáš Nagy, and Michael Pinsker. An order out of nowhere: A new algorithm for infinite-domain CSPs. In *Proceedings of the 51st International Colloquium on Automata, Languages, and Programming - ICALP'24*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 148:1–148:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ICALP.2024.148.
- 60 Antoine Mottet and Michael Pinsker. Smooth approximations: An algebraic approach to CSPs over finitely bounded homogeneous structures. *J. ACM.*, 71(5):36:1–36:47, 2024. doi:10.1145/3689207.
- 61 Jaroslav Opatrny. Total ordering problem. *SIAM J. Comput.*, 8(1):111–114, 1979. doi:10.1137/0208008.
- 62 Michael Pinsker. Current challenges in infinite-domain constraint satisfaction: Dilemmas of the infinite sheep. In *Proceedings of the 52nd IEEE International Symposium on Multiple-Valued Logic, ISMVL'22*, pages 80–87. IEEE, 2022. doi:10.1109/ISMVL52857.2022.00019.
- 63 Michael Pinsker and Clemens Schindler. On the Zariski topology on endomorphism monoids of omega-categorical structures. *J. Symb. Log.*, 2023. To appear. doi:10.1017/jsl.2023.81.
- 64 J. Rydval. Finitely Bounded Homogeneity Turned Inside-Out. To appear in *J. Math. Log.*, preprint: arXiv.2108.00452, 2025.
- 65 Lynn Scow. Ramsey transfer to semi-retractions. *Ann. Pure Appl. Log.*, 172(3):Article 102891, 2021. doi:10.1016/j.apal.2020.102891.
- 66 Mark H. Siggers. A strong Mal'cev condition for locally finite varieties omitting the unary type. *Algebra Univers.*, 64:15–20, 2010. doi:10.1007/s00012-010-0082-3.
- 67 Paul Winkler. *On the Descriptive Complexity of Phylogeny Constraint Satisfaction Problems*. Master's thesis, Technische Universität Wien, 2024. available online.
- 68 Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science - FOCS'17*, pages 331–342. IEEE, 2017. doi:10.1109/FOCS.2017.38.
- 69 Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM.*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.