COMPLEXITY CLASSIFICATION TRANSFER FOR CSPS VIA ALGEBRAIC PRODUCTS *

MANUEL BODIRSKY†, PETER JONSSON‡, BARNABY MARTIN§, ANTOINE MOTTET¶, AND ŽANETA SEMANIŠINOVÁ†

Abstract. We study the complexity of infinite-domain constraint satisfaction problems: our basic setting is that a complexity classification for the CSPs of first-order expansions of a structure $\mathfrak{A}$ can be transferred to a classification of the CSPs of first-order expansions of another structure $\mathfrak{B}$. We exploit a product of structures (the *algebraic product*) that corresponds to the product of the respective polymorphism clones and present a complete complexity classification of the CSPs for first-order expansions of the n-fold algebraic power of $(\mathbb{Q};<)$. This is proved by various algebraic and logical methods in combination with knowledge of the polymorphisms of the tractable first-order expansions of $(\mathbb{Q};<)$ and explicit descriptions of the expressible relations in terms of syntactically restricted first-order formulas. By combining our classification result with general classification transfer techniques, we obtain surprisingly strong new classification results for highly relevant formalisms such as Allen's Interval Algebra, the n-dimensional Block Algebra, and the Cardinal Direction Calculus, even if higher-arity relations are allowed. Our results confirm the infinite-domain tractability conjecture for classes of structures that have been difficult to analyse with older methods. For the special case of structures with binary signatures, the results can be substantially strengthened and tightly connected to Ord-Horn formulas; this solves several longstanding open problems from the AI literature.

Key words. constraint satisfaction, temporal reasoning, computational complexity, polymorphisms, universal algebra, polynomial-time tractability

MSC codes. 06A05, 68Q25, 08A70

1. Introduction. This introductory section is divided into three parts where we describe the background, present our contributions, and provide an outline of the article, respectively.

Background. Constraint satisfaction problems (CSPs) are computational problems that appear in many areas of computer science, for example in temporal and spatial reasoning in artificial intelligence [26] or in database theory [69, 18]. The computational complexity of CSPs is of central interest in these areas, and a general research goal is to obtain systematic complexity classification results, in particular about CSPs that are in P and CSPs that are NP-hard. CSPs can be described elegantly by fixing a structure with a finite relational signature, the *template*; the computational task is to determine whether a given finite input structure has a homomorphism to the template. A breakthrough result was obtained independently by Bulatov [46] and by Zhuk [100, 99], which confirmed the famous Feder-Vardi conjecture [58]: every CSP over a finite template (i.e., a structure with a finite domain) is in P, or it is NP-complete. Moreover, given the template it is possible to decide algorithmically whether its CSP is in P or whether it is NP-complete.

Most of the CSPs in temporal and spatial reasoning can *not* be formulated as CSPs with a finite

*Manuel Bodirsky and Žaneta Semanišinová have received funding from the ERC (Grant Agreement no. 101071674, POCOCOP) and from the DFG (Project FinHom, Grant 467967530). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Peter Jonsson is partially supported by the Swedish Research Council (VR) under grants 2017-04112 and 2021-04371. Barnaby Martin is supported by EPSRC grant EP/X03190X/1.

†TU Dresden, Germany, (manuel.bodirsky@tu-dresden.de, zaneta.semanisinova@tu-dresden.de)

‡Linköping University, Sweden, (peter.jonsson@liu.se)

§Durham University, UK, (barnaby.d.martin@durham.ac.uk)

¶Hamburg University of Technology, Research Group for Theoretical Computer Science, Germany, (antoine.mottet@tuhh.de)

template. The same is true for many of the CSPs that appear in database theory (e.g., most of the CSPs in the logic MMSNP, which is a fragment of existential second-order logic introduced by Feder and Vardi [58], and which is important for database theory [18], cannot be formulated as CSPs with a finite template [78]). For CSPs with infinite templates we may not hope for general classification results [25]; however, we may hope for general classification results if we restrict our attention to classes of templates that are model-theoretically well behaved. An example of such a class is the class of all structures with domain $\mathbb{Q}$ where all relations are definable with a first-order formula over the structure $(\mathbb{Q}; <)$. This class of structures is of fundamental interest in model theory, and, by a result of Cameron [48], also in the theory of infinite permutation groups (it is precisely the class of all countable structures with a highly set-transitive automorphism group). The CSPs for such structures have been called *temporal CSPs* because they include many CSPs that are of relevance in temporal reasoning, such as the Betweenness problem [88], the And/Or scheduling problem [80], or the satisfiability problem for Ord-Horn constraints [87]. The complexity of temporal CSPs has been classified by Bodirsky and Kára [30], and a temporal CSP is either in P or it is NP-complete.

Over the past 10 years, many classes of infinite structures have been classified with respect to the complexity of their CSP. We may divide these results into *first-* and *second-generation* classifications. First-generation classifications, such as the classification of temporal CSPs mentioned above, typically use concepts from universal algebra and Ramsey theory, and essentially proceed by a combinatorial case distinction [41, 35, 70, 28]. Second-generation classifications also use universal algebra and Ramsey theory, but they eliminate large parts of the combinatorial analysis by using arguments for finite structures and the Bulatov-Zhuk theorem (the first example following this approach is [36]; other examples include [34, 84, 20, 82, 33]). So the idea of second-generation classifications is to *transfer* the finite-domain classification to certain tame classes of infinite structures.

Complexity transfer is also the topic of the present article; however, we transfer classification results not from finite structures to classes of infinite structures, but between classes of infinite structures. The key to systematically relating many classes of infinite structures are various *product constructions*, and *logical interpretations*. Examples are *Allen's Interval Algebra* [3] from temporal reasoning, which has a first-order interpretation in $(\mathbb{Q}; <)$, or the *rectangle algebra* [63, 86] and the r-dimensional *block algebra* [10]. These links extend to links between fragments of the respective formalisms. In order to also establish links between the complexity of the respective CSPs, the logical interpretations must use *primitive positive* formulas, rather than full first-order logic. There are various notions of products of constraint formalisms that have been studied in the literature; see [96, 47]. In this article we use a product of structures known as the *algebraic product*; it corresponds to the product of the respective polymorphism clones, which is essential for the universal algebraic approach.

Contributions. This article contains both theoretical results and applications of these results to well-studied formalisms and open problems in the area. Our first main contribution is a complete complexity classification for the CSPs of first-order expansions of $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$, i.e., the algebraic product of $(\mathbb{Q}; <)$ with itself. We then generalise this result to first-order expansions of finite algebraic powers of $(\mathbb{Q}; <)$, denoted by $(\mathbb{Q}; <)^{(n)}$. In the proof we use known results about first-order expansions of $(\mathbb{Q}; <)$ combined with a mix of algebraic and of logical arguments. On the algebraic side, we use the fact that the first-order expansions of $(\mathbb{Q}; <)$ with a tractable CSP have certain polymorphisms. On the logic side, we use highly informative descriptions of the relations of the templates using syntactically restricted forms of first-order logic, so-called (weakly) *i-determined clauses*. These syntactic forms partially allow us to separate the relations coming from different factors of the algebraic product. The combination of algebraic and syntactic methods turned out to be very powerful in our setting and we believe that it will be fruitful for analysing first-order expansions of products of other structures.

Together with a general classification transfer result from [19], we then obtain a sequence of new complexity classification results for classes of CSPs that have been studied in temporal and spatial reasoning. We derive our applications in two steps: we first derive classification results for structures with relations of arbitrary arity. With little extra effort, we then obtain stronger results for the special case that all relations are binary. The restriction to binary relations is very common in the AI and qualitative reasoning literature: the influential survey by Dylla et al. [56] lists 50+ formalisms based on binary relations but just a handful of formalisms that use relations with higher arity. One reason for this is that binary relations in the infinite-domain regime are very powerful and have attracted much attention. Another reason is pointed out in [56, Section 3]: generalizations to higher-arity relations are indeed useful in practice, but progress has been hampered by a lack of algebraic understanding and algorithmic methods. The results in this article provide a step towards a better understanding of AI-relevant formalisms and their generalisation to non-binary relations.

Templates with relations of unrestricted arity. We determine the complexity of the CSP for first-order expansions of the basic relations in three influential formalisms for spatio-temporal reasoning: Allen's Interval Algebra [3], the Block Algebra (BA) [10], and the Cardinal Direction Calculus (CDC) [76]. Allen's interval algebra is a fundamental formalism within AI and qualitative reasoning that has a myriad of applications, e.g. in automated planning [5, 85, 89], natural language processing [53, 95] and molecular biology [61]. Both BA and CDC can be viewed as variations of Allen's algebra originally aimed at expanding the range of applicability. For instance, the BA can handle directions in a spatial reasoning setting (something that is difficult in standard formalisms such as RCC [91]) with diverse applications such as computer vision [51], architecture [92], and physics simulation in computer games [98]. CDC have found applications in geographic information systems and image interpretation: see [49, 59, 76] and the references therein. These three formalisms have been very important in the evolution of calculi for qualitative reasoning, for the development of methods for complexity classifications, and for the algebraic theory of infinite-domain CSPs. For instance, two milestones in complexity classification of infinite-domain CSPs are concerned with Allen's algebra: Nebel and Bürckert's [87] result for subsets of binary Allen relations containing all basic relations, and Krokhin et al's [72] generalization to arbitrary subsets of binary Allen relations. BA and CDC, on the other hand, have presented significant challenges and resisted full complexity classifications for at least twenty years. We conclude that Allen's Interval Algebra, BA and CDC were important test cases for CSP complexity classification projects long before any complexity classification conjectures for infinite-domain CSPs had been formulated.

In these particular cases, our results show that the so-called *infinite-domain tractability conjecture*, which is formulated for all reducts of finitely bounded homogeneous structures [44], holds. The conjecture states that such a structure has a polynomial-time tractable CSP unless the structure admits a primitive positive interpretation of a structure which is homomorphically equivalent to K_3, the clique with three vertices (note that the CSP for the template K_3 is the 3-colourability problem, which is a well-known NP-complete problem). This hardness condition is known to be equivalent to the structure admitting a primitive positive construction of K_3 [13]. All the classes of infinite structures discussed so far are first-order interpretable over $(\mathbb{Q}; <)$ and it can be shown that they fall into the scope of this conjecture (see, for instance, Theorem 4 from [83], Lemma 3.5.4 and Proposition 4.2.19 from [19], and Lemma 3.8 from [13]). Interestingly, the structures we treat here are notoriously difficult for the methods underpinning second-generation classification results:

e.g., the unique interpolation property usually fails in this context [20].

To make progress with proving the infinite-domain tractability conjecture, one strategy is to verify it on larger and larger classes of structures. Highly useful restrictions on classes of interesting structures come from model theory. The concept of *stability* and, more generally, *NIP* (i.e., not having Shelah's *independence property* [93]) are central concepts in model theory (see, e.g., [94, 50]). While the consequences of these concepts for the complexity of constraint satisfaction are unclear, stability and NIP are still relevant here, because in combination with homogeneity in a finite relational signature they have strong consequences and allow for model-theoretic classification results, which in turn can be the basis for CSP classification results.

We would like to stress the particular role of structures with a first-order interpretation over $(\mathbb{Q};<)$ in this context. All of these structures are NIP. Moreover, it is known that every homogeneous structure with a finite relational signature which is stable has a first-order interpretation in $(\mathbb{Q};<)$ [73]. These structures are not only important in model theory, but also significant algorithmically, because the cases with a polynomial-time tractable CSP usually cannot be characterised by canonical polymorphisms, and new algorithms rather than polynomial-time reductions to tractable finite-domain CSPs are needed [31, 40, 84]. Therefore, a complexity classification for CSPs of structures with a first-order interpretation in $(\mathbb{Q};<)$ would be an important milestone for resolving the tractability conjecture. Since our syntactic approach to complexity classification overcomes the mentioned challenges in important cases, it represents a step towards this goal.

Templates with binary relations. Our results concerning first-order expansions of $(\mathbb{Q};<)^{(n)}$ can be specialised to the case when only binary relations are allowed. If $\mathfrak{D}$ is such a structure, then our results imply that $\mathrm{CSP}(\mathfrak{D})$ is in P if and only if every relation in $\mathfrak{D}$ can be defined by an Ord-Horn formula [87]. This allows us to answer several open questions from the AI literature. In particular, we solve an open problem from 2002 about the n-dimensional cardinal direction calculus [8] (Section 6.1), an open problem from 1999 about fragments of the rectangle algebra [9] (Theorem 6.8) and an open problem from 2002 about the n-dimensional block algebra [10] (Corollary 6.10). We can also answer another question in [10] about integration of the tractable cases into tractable formalisms that can also handle metric constraints; see the discussion at the end of Section 6.3. Finally, we obtain short new proofs of known results about reducts of Allen's Interval Algebra (Section 6.2). Our results typically answer more general questions than those asked in the publications above, in particular, they yield results also in the case when the relations are of arity higher than two.

Outline. The structure of the article is as follows. Section 2 contains the basic concepts that are needed for a formal definition of the CSPs, and some facts about constraint satisfaction problems and their computational complexity. Section 3 contains the definition of the algebraic product together with some related results. In Section 4, we study $(\mathbb{Q};<)^{(n)}$, and this ultimately provides us with a complexity classification of the CSP for first-order expansions of $(\mathbb{Q};<)^{(n)}$. We additionally study the restriction to binary signatures in this section, i.e., signatures where all relations have arity at most two. The next section is devoted to a condensed introduction to complexity classification transfer. Thereafter, we combine the complexity results for $(\mathbb{Q};<)^{(n)}$ with complexity classification transfer in order to analyse various spatio-temporal formalisms in Section 6. We conclude the article with a brief discussion of the results together with some possible future research directions (Section 7).

Some of the results in this article have been announced in a conference paper [27]. However, one of the central proofs there (Lemma 2) is not correct. The proof in the present article avoids proving Lemma 2 from [27] and is entirely new; in particular, the syntactic approach to analysing first-order expansions of products in Section 4.3 did not appear in the old approach.

2. Constraint Satisfaction Problems. In this section we introduce basic concepts that are needed for a formal definition of the class of constraint satisfaction problems (CSP) together with some basic facts about CSPs and their computational complexity.

2.1. Basic Definitions. Let τ be a *relational signature*, i.e., a set of *relation symbols* R, each equipped with an *arity* $k \in \mathbb{N}$. A *τ-structure* $\mathfrak{A}$ consists of a set A, called the *domain* of $\mathfrak{A}$, and a relation $R^{\mathfrak{A}} \subseteq A^k$ for each relation symbol $R \in \tau$ of arity k. A structure is called *finite* if its domain is finite. Relational structures are often written like $(A; R_1^{\mathfrak{A}}, R_2^{\mathfrak{A}}, \dots)$, with the obvious interpretation; for example, $(\mathbb{Q}; <)$ denotes the structure whose domain is the set of rational numbers $\mathbb{Q}$ and which carries a single binary relation $<$ which denotes the usual strict order of the rationals. Sometimes, we do not distinguish between the symbol R for a relation and the relation $R^{\mathfrak{A}}$ itself. 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 $R^{\mathfrak{A}} = R^{\mathfrak{A}'}$ for all $R \in \tau$, then $\mathfrak{A}$ is called a *τ-reduct* (or simply *reduct*) of $\mathfrak{A}'$, and $\mathfrak{A}'$ is called a *τ'-expansion* (or simply *expansion*) of $\mathfrak{A}$. If R is a relation over the domain of $\mathfrak{B}$, then we let $(\mathfrak{A}; R)$ denote the expansion of $\mathfrak{A}$ by R.

We continue by introducing some logical terminology and machinery. We refer the reader to [65] for an introduction to first-order logic. An *atomic τ-formula* is a formula of the form $x = y$, $R(x_1, \dots, x_n)$, or the form $\bot$, where $x_1, \dots, x_n, x, y$ are variables, R is a symbol from τ, and $\bot$ is a symbol that stands for 'false'. Let $\mathfrak{B}$ denote a τ-structure. If ψ is a sentence (i.e. a first-order formula without free variables), then we write $\mathfrak{B} \models \psi$ to denote that $\mathfrak{B}$ is a model of (or satisfies) ψ. One can use first-order formulas over the signature τ to define relations over $\mathfrak{B}$: if $\phi(x_1, \dots, x_n)$ is a first-order τ-formula with free variables $x_1, \dots, x_n$, then the relation *defined* by ϕ over $\mathfrak{B}$ is the relation $\{(b_1, \dots, b_n) \in B^n \mid \mathfrak{B} \models \phi(b_1, \dots, b_n)\}$. We say that τ-formulas $\phi(x_1, \dots, x_n)$ and $\psi(x_1, \dots, x_n)$ are equivalent over $\mathfrak{B}$, if $\mathfrak{B} \models \forall x_1, \dots, x_n (\phi \Leftrightarrow \psi)$. We often omit the specification of the structure if it is clear from the context. We say that a structure $\mathfrak{B}$ has *quantifier elimination* if every first-order formula is equivalent to a quantifier-free formula over $\mathfrak{B}$. Every quantifier-free formula can be written in *conjunctive normal form* (CNF), i.e., as a conjunction of disjunctions of *literals*, i.e., atomic formulas or their negations. A disjunction of literals is also called a *clause*.

If $\mathfrak{A}$ and $\mathfrak{B}$ are τ-structures, then a *homomorphism* from $\mathfrak{A}$ to $\mathfrak{B}$ is a function $h \colon A \to B$ that *preserves* all the relations, that is, if $(a_1, \dots, a_k) \in R^{\mathfrak{A}}$, then $(h(a_1), \dots, h(a_k)) \in R^{\mathfrak{B}}$. The structures $\mathfrak{A}$ and $\mathfrak{B}$ are called *homomorphically equivalent* if there exists a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$ and a homomorphism from $\mathfrak{B}$ to $\mathfrak{A}$. A first-order τ-formula is *preserved* by a map between two τ-structures $\mathfrak{A}$ and $\mathfrak{B}$ if it preserves the relation defined by the formula in these structures.

A *first-order expansion* of $\mathfrak{A}$ is a structure $\mathfrak{A}'$ augmented by relations that are first-order definable in $\mathfrak{A}$. A *first-order reduct* of $\mathfrak{A}$ is a reduct of a first-order expansion of $\mathfrak{B}$. Relational structures might have an infinite signature; however, to avoid representational issues and for simplicity we restrict ourselves to finite signatures in the following definition.

Definition 2.1 (CSPs). *Let τ be a finite relational signature and let $\mathfrak{B}$ be a τ-structure. The* constraint satisfaction problem *for $\mathfrak{B}$, denoted by* $\mathrm{CSP}(\mathfrak{B})$*, is the computational problem of deciding for a given finite τ-structure $\mathfrak{A}$ whether $\mathfrak{A}$ has a homomorphism to $\mathfrak{B}$ or not.*

Note that this definition of constraint satisfaction problems can be used even if $\mathfrak{B}$ is an infinite structure over a finite relational signature. Also note that homomorphically equivalent structures have the same CSP.

Example 2.2. The structure $(\{0,1,2\};\neq)$ is denoted by K_3. The problem $\mathrm{CSP}(K_3)$ is the three-colorability problem for graphs. The input is a structure with a single binary relation, representing edges in a graph (ignoring the orientation); homomorphisms from this graph to K_3 correspond precisely to the proper 3-colorings of the graph.

2.2. Primitive Positive Constructions. Three central concepts in the complexity analysis of CSPs are *primitive positive definitions*, *primitive positive interpretations*, and *primitive positive constructions*. The three concepts are increasingly powerful. Their definitions build on each other and will be recalled here for the convenience of the reader.

A *primitive positive τ-formula* is a formula $\phi(x_1,\dots,x_n)$ with free variables $x_1,\dots,x_n$ of the form

$$\exists y_1,\dots,y_l(\psi_1 \wedge \dots \wedge \psi_m)$$

where $\psi_1,\dots,\psi_k$ are atomic τ-formulas over the variables $x_1,\dots,x_n,y_1,\dots,y_l$. Two relational structures $\mathfrak{A}$ and $\mathfrak{B}$ are called

- *(primitively positive) interdefinable* if they have the same domain $A = B$, and if every relation of $\mathfrak{A}$ is (primitively positively) definable in $\mathfrak{B}$ and vice versa.
- *(primitively positively) bi-definable* if $\mathfrak{B}$ is isomorphic to a structure that is (primitively positively) interdefinable with $\mathfrak{A}$.

We will now turn our attention towards methods for complexity analysis.

LEMMA 2.3 ([66]). *Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures with finite relational signatures and the same domain. If every relation of $\mathfrak{A}$ has a primitive positive definition in $\mathfrak{B}$, then there is a polynomial-time reduction from* $\mathrm{CSP}(\mathfrak{A})$ *to* $\mathrm{CSP}(\mathfrak{B})$.

Primitive positive definability can be generalised as follows.

DEFINITION 2.4 (Interpretations). *A* (primitive positive) interpretation *of a structure $\mathfrak{C}$ in a structure $\mathfrak{B}$ is a partial surjection I from B^d to C, for some finite $d \in \mathbb{N}$ called the* dimension *of the interpretation, such that the preimage of a relation of arity k defined by an atomic formula in $\mathfrak{C}$, considered as a relation of arity dk over B, is (primitively positively) definable in $\mathfrak{B}$; in this case, we say that $\mathfrak{C}$ is (primitively positively) interpretable in $\mathfrak{B}$.*

Two structures $\mathfrak{B}$ and $\mathfrak{C}$ such that $\mathfrak{B}$ has a primitive positive interpretation in $\mathfrak{C}$ and $\mathfrak{C}$ has a primitive positive interpretation in $\mathfrak{B}$ are called mutually primitively positively interpretable.

Note that in particular $x = x$ and $x = y$ are atomic formulas and hence the domain and the kernel of I are primitively positively definable. Primitive positive interpretations preserve the complexity of CSPs in the following way.

PROPOSITION 2.5 (see, e.g., Theorem 3.1.4 in [19]). *Let $\mathfrak{B}$ and $\mathfrak{C}$ be structures with finite relational signatures. If $\mathfrak{C}$ has a primitive positive interpretation in $\mathfrak{B}$, then there is a polynomial-time reduction from* $\mathrm{CSP}(\mathfrak{C})$ *to* $\mathrm{CSP}(\mathfrak{B})$.

Example 2.6. Let $\mathbb{I}$ be the set of all pairs $(x,y) \in \mathbb{Q}^2$ such that $x < y$; i.e., $\mathbb{I}$ might be viewed as the set of all closed intervals $[a,b]$ of rational numbers. Let m be the binary relation over $\mathbb{I}$ that contains all pairs $((a_1,a_2),(b_1,b_2))$ such that $a_2 = b_1$. Then the structure $(\mathbb{I};\mathsf{m})$ has a primitive positive interpretation (of dimension 2) in $(\mathbb{Q};<)$:

- The interpretation map is the identity map on $\mathbb{I} \subseteq \mathbb{Q}^2$.
- The preimage of the relation defined by the atomic formula $a = b$ is defined by the formula $a_1 = b_1 \wedge a_2 = b_2 \wedge a_1 < a_2$.

TABLE 2.1
Basic relations in the interval algebra.

Basic relation		Example	Endpoints
X precedes Y	p	XXX	$X^+ < Y^-$
Y preceded by X	$\mathsf{p}^{\smile}$	YYY	
X meets Y	m	XXXX	$X^+ = Y^-$
Y is met by X	$\mathsf{m}^{\smile}$	YYYY	
X overlaps Y	o	XXXX	$X^- < Y^- \wedge$
		YYYY	$Y^- < X^+ \wedge$
Y overlapped by X	$\mathsf{o}^{\smile}$		$X^+ < Y^+ \wedge$
X during Y	d	XX	$X^- > Y^- \wedge$
Y includes X	$\mathsf{d}^{\smile}$	YYYYYY	$X^+ < Y^+$
X starts Y	s	XXX	$X^- = Y^- \wedge$
Y started by X	$\mathsf{s}^{\smile}$	YYYYYY	$X^+ < Y^+$
X finishes Y	f	XXX	$X^+ = Y^+ \wedge$
Y finished by X	$\mathsf{f}^{\smile}$	YYYYYY	$X^- > Y^-$
X equals Y	$\equiv$	XXXX	$X^- = Y^- \wedge$
		YYYY	$X^+ = Y^+$

- The preimage of the relation defined by the atomic formula $\mathsf{m}(a, b)$ is defined by the formula $a_1 < a_2 \wedge a_2 = b_1 \wedge b_1 < b_2$;

It is straightforward to adapt the construction above to atomic formulas that are obtained by variable identification by using the equality relation. Proposition 2.5 implies that $\mathrm{CSP}(\mathbb{I}; \mathsf{m})$ is in P since $\mathrm{CSP}(\mathbb{Q}; <)$ is in P.

Example 2.7. The *interval algebra* [3] is a formalism that is both well-known and well-studied in AI. It can be viewed as a relational structure with the domain $\mathbb{I}$ introduced in Example 2.6 and a binary relation symbol for each binary relation $R \subseteq \mathbb{I}^2$ such that the relation $\{(a_1, a_2, b_1, b_2) \mid ((a_1, a_2), (b_1, b_2)) \in R\}$ is first-order definable in $(\mathbb{Q}; <)$. We let $\mathfrak{IA}$ denote this structure and we let $\top$ denote the relation which holds for all pairs of intervals. Clearly, Allen's Interval Algebra has a 2-dimensional interpretation in $(\mathbb{Q}; <)$, but not a primitive positive interpretation.

The *basic* relations of Allen's Interval Algebra are the 13 relations defined in Table 2.1: we let $\mathfrak{IA}^b$ be the corresponding structure. If $I = [a, b] \in \mathbb{I}$, then we write I^- for a and I^+ for b. It is well-known that all the basic relations of Allen's Interval Algebra have a primitive positive definition over $(\mathbb{I}; \mathsf{m})$ [4]. We conclude that $\mathrm{CSP}(\mathfrak{IA}^b)$ is in P since $\mathrm{CSP}(\mathbb{I}; \mathsf{m})$ is in P by the previous example.

We finally define primitive positive constructions. We will not use such constructions in our proofs but it is used in the statement of Theorem 2.11 and, thus, in the formulation of the infinite-domain tractability conjecture.

DEFINITION 2.8 (Primitive positive constructions). *A structure* $\mathfrak{C}$ *has a* primitive positive construction *in* $\mathfrak{B}$ *if* $\mathfrak{C}$ *is homomorphically equivalent to a structure* $\mathfrak{C}'$ *with a primitive positive interpretation in* $\mathfrak{B}$.

LEMMA 2.9. *Let* $\mathfrak{B}$ *and* $\mathfrak{C}$ *be structures with finite relational signature. If* $\mathfrak{C}$ *has a primitive positive construction in* $\mathfrak{B}$*, then there is a polynomial-time reduction from* $\mathrm{CSP}(\mathfrak{C})$ *to* $\mathrm{CSP}(\mathfrak{B})$.

Proof. An immediate consequence of Proposition 2.5 since homomorphically equivalent structures have the same CSP. ☐

2.3. Model Theory and Algebra. This section collects some basic terminology and facts from model theory and algebra. The set of all first-order τ-sentences that are true in a given τ-structure $\mathfrak{A}$ is called the *first-order theory* of $\mathfrak{A}$. A countable structure $\mathfrak{A}$ is *ω-categorical* if all countable models of the first-order theory of $\mathfrak{A}$ are isomorphic. The structure $(\mathbb{Q}; <)$, and all structures with an interpretation in $(\mathbb{Q}; <)$, are ω-categorical: for the structure $(\mathbb{Q}; <)$, this was shown by Cantor who proved that all countable dense and unbounded linear orders are isomorphic. All structures that we encounter in the later parts of this article are ω-categorical.

An *automorphism* of a structure $\mathfrak{A}$ is a permutation α of A such that both α and its inverse are homomorphisms. The set of all automorphisms of a structure $\mathfrak{A}$ is denoted by $\mathrm{Aut}(\mathfrak{A})$, and forms a group with respect to composition; the neutral element of the group is the identity map which is denoted by id_A. The set of all permutations of a set A is called the *symmetric group* and denoted by $\mathrm{Sym}(A)$. Clearly, $\mathrm{Sym}(A)$ equals $\mathrm{Aut}(A; =)$.

The *orbit* of $(a_1, \dots, a_n) \in A^n$ in $\mathrm{Aut}(\mathfrak{A})$ is the set $\{(\alpha(a_1), \dots, \alpha(a_n)) \mid \alpha \in \mathrm{Aut}(\mathfrak{A})\}$. It was proved independently by Engeler, Ryll-Nardzewski, and Svenonius that a countable structure $\mathfrak{A}$ is ω-categorical if and only if $\mathrm{Aut}(\mathfrak{A})$ is *oligomorphic*, i.e., has only finitely many orbits of n-tuples, for all $n \geq 1$ (see, e.g., [65, Theorem 6.3.1]). This implies that structures with a first-order interpretation in an ω-categorical structure are ω-categorical [65, Theorem 6.3.6]. In particular, first-order reducts of ω-categorical structures are again ω-categorical. It also follows that first-order expansions of ω-categorical structures $\mathfrak{B}$ are ω-categorical themselves since such relations are preserved by all automorphisms of $\mathfrak{B}$. In an ω-categorical structure $\mathfrak{B}$, a relation is preserved by all automorphisms of $\mathfrak{B}$ if and only if it is first-order definable in $\mathfrak{B}$ (see, e.g., [19, Theorem 4.2.9]).

Another important property of $(\mathbb{Q}; <)$ is called *homogeneity*. An *embedding* from $\mathfrak{A}$ to $\mathfrak{B}$ is an injective homomorphism e from $\mathfrak{A}$ to $\mathfrak{B}$ which also preserves the complement of each relation, i.e., $(a_1, \dots, a_k) \in R^{\mathfrak{A}}$ if and only if $(f(a_1), \dots, f(a_k)) \in R^{\mathfrak{B}}$. A τ-structure $\mathfrak{A}$ is called a *substructure* of $\mathfrak{B}$ if $A \subseteq B$ and the identity map id_A is an embedding from $\mathfrak{A}$ to $\mathfrak{B}$. Now, a relational structure is called *homogeneous* (or sometimes *ultrahomogeneous*) if every isomorphism between finite substructures can be extended to an automorphism of the structure [65, p. 160]. An ω-categorical structure $\mathfrak{B}$ is homogeneous if and only if $\mathfrak{B}$ has quantifier elimination [65, Theorem 6.4.1].

As mentioned above, a relation R is first-order definable in an ω-categorical structure $\mathfrak{B}$ if and only if it is preserved by all automorphisms of $\mathfrak{B}$. Similarly, the question whether a given relation is primitively positively definable in $\mathfrak{B}$ can be studied using *polymorphisms*. A polymorphism of a structure $\mathfrak{B}$ is a homomorphism from $\mathfrak{B}^k$ to $\mathfrak{B}$. Here, the structure $\mathfrak{B}^k$ denotes the k-fold direct product structure $\mathfrak{B} \times \dots \times \mathfrak{B}$; more generally, if $\mathfrak{B}_1, \dots, \mathfrak{B}_k$ are τ-structures, then $\mathfrak{C} = \mathfrak{B}_1 \times \dots \times \mathfrak{B}_k$ is defined to be the τ-structure with domain $B_1 \times \dots \times B_k$ and for every $R \in \tau$ of arity m we have

$$\begin{aligned} R^{\mathfrak{C}} = \{ & ((a_{1,1}, \dots, a_{1,k}), \dots, (a_{m,1}, \dots, a_{m,k})) \in C^m \mid \\ & (a_{1,1}, \dots, a_{m,1}) \in R^{\mathfrak{B}_1}, \dots, (a_{1,k}, \dots, a_{m,k}) \in R^{\mathfrak{B}_k} \}. \end{aligned}$$

The set of all polymorphisms of a structure $\mathfrak{B}$ is denoted by $\mathrm{Pol}(\mathfrak{B})$. *Endomorphisms* are a special case of polymorphisms with $k = 1$: an endomorphism of a structure $\mathfrak{B}$ is thus a homomorphism from $\mathfrak{B}$ to itself. The set of all endomorphisms of $\mathfrak{B}$ is denoted by $\mathrm{End}(\mathfrak{B})$. For every $i \leq n$, the i-th projection of arity n is the operation π_i^n defined by $\pi_i^n(x_1, \dots, x_n) \coloneqq x_i$. The set of all polymorphisms of a structure $\mathfrak{B}$ forms an *(operation) clone*: it is closed under composition and contains all projections. Moreover, an operation clone $\mathscr{C}$ on a set B is a polymorphism clone of a relational structure if and only if the operation clone is *closed*, i.e., for each $k \geq 1$ the set of k-ary operations in $\mathscr{C}$ is closed with respect to the product topology on B^{B^k} where B is taken to be discrete (see, e.g., Corollary 6.1.6 in [19]).

An operation clone $\mathscr{C}$ is called *oligomorphic* if the permutation group $\mathscr{G}$ of invertible unary maps in $\mathscr{C}$ is oligomorphic. A relation $R \subseteq B^n$ is preserved by all polymorphisms of an ω-categorical structure $\mathfrak{B}$ if and only if R has a primitive positive definition in $\mathfrak{B}$ [38]. It follows that two ω-categorical relational structures with the same domain have the same polymorphisms if and only if they are primitively positively interdefinable. If $\mathscr{F} \subseteq \mathscr{C}$, we write

- $\langle \mathscr{F} \rangle$ for the smallest subclone of $\mathscr{C}$ which contains $\mathscr{F}$, and
- $\overline{\mathscr{F}}$ for the smallest *(locally, or topologically) closed* subset of $\mathscr{C}$ that contains $\mathscr{F}$: that is, $\overline{\mathscr{F}}$ consists of all operations f such that for every finite subset S of the domain there exists an operation $g \in \mathscr{F}$ which agrees with f on S.

If $\sigma\colon \{1,\dots,n\} \to \{1,\dots,k\}$ and $f \in \mathscr{C}$ has arity n, then we write f^σ for the operation $f(\pi^k_{\sigma(1)},\dots,\pi^k_{\sigma(n)})$ which maps $(x_1,\dots,x_k)$ to $f(x_{\sigma(1)},\dots,x_{\sigma(n)})$. Let $\mathscr{C}_1$ and $\mathscr{C}_2$ be two clones and let ξ be a function from $\mathscr{C}_1$ to $\mathscr{C}_2$ that preserves the arities of the operations. Then ξ is

- a *clone homomorphism* if for all $f \in \mathscr{C}_1$ of arity n and $g_1,\dots,g_n \in \mathscr{C}_1$ of arity k we have
$$\xi(f(g_1,\dots,g_n)) = \xi(f)(\xi(g_1),\dots,\xi(g_n))$$
and $\xi(\pi^n_i) = \pi^n_i$ for all $i \leq n$.
- *minor-preserving* if for all $f \in \mathscr{C}_1$ of arity n and $\sigma\colon \{1,\dots,n\} \to \{1,\dots,k\}$
$$\xi(f^\sigma) = \xi(f)^\sigma.$$
- *uniformly continuous* if for every finite subset F of the domain of $\mathscr{C}_2$ there exists a finite subset G of the domain of $\mathscr{C}_1$ such that and for every $n \geq 1$ and all $f, g \in \mathscr{C}_1$ of arity n, if $f|_{G^n} = g|_{G^n}$, then $\xi(f)|_{F^n} = \xi(g)|_{F^n}$.

Note that trivially, every clone homomorphism is minor-preserving. Uniformly continuous clone homomorphims naturally arise from interpretations.

LEMMA 2.10 ([42]). *If $\mathfrak{C}$ has a primitive positive interpretation in $\mathfrak{B}$, then* $\mathrm{Pol}(\mathfrak{B})$ *has a uniformly continuous clone homomorphism to* $\mathrm{Pol}(\mathfrak{C})$*.*

The relevance of minor-preserving maps for the complexity of constraint satisfaction is witnessed by the following theorem.

THEOREM 2.11 ([13]). *If $\mathfrak{B}$ is an ω-categorical structure and $\mathfrak{C}$ is a finite structure, then $\mathfrak{C}$ has a primitive positive construction in $\mathfrak{B}$ if and only if* $\mathrm{Pol}(\mathfrak{B})$ *has a uniformly continuous minor-preserving map to* $\mathrm{Pol}(\mathfrak{C})$*.*

Note that Theorem 2.11 in combination with Lemma 2.9 implies the following.

COROLLARY 2.12. *Let $\mathfrak{B}$ be an ω-categorical structure and suppose that* $\mathrm{Pol}(\mathfrak{B})$ *has a uniformly continuous minor-preserving map to* $\mathrm{Pol}(K_3)$*. Then $\mathfrak{B}$ has a finite-signature reduct whose* CSP *is NP-hard.*

The *infinite-domain tractability conjecture* implies that for reducts of finitely bounded homogeneous structures, and if P $\neq$ NP, then the condition given in Corollary 2.12 is not only sufficient, but also necessary for NP-hardness [44]. Note that CSPs of reducts of finitely bounded homogeneous structures are always in NP (see, e.g., Proposition 2.3.15 in [19]). The conjecture has also interesting algebraic interpretations in line with the dichotomy for CSPs of finite structures [14, 12, 13].

An operation $f\colon B^k \to B$ is called a *weak near unanimity operation* if for all $x, y \in B$ the operation f satisfies

$$f(y, x, \dots, x) = f(x, y, \dots, x) = \cdots = f(x, \dots, x, y).$$

If $\mathscr{C}$ is a clone on a finite domain B without a minor-preserving map to $\mathrm{Pol}(K_3)$, then $\mathscr{C}$ contains a weak near unanimity operation [79]. A potential generalisation of this fact to polymorphism clones of ω-categorical structures $\mathfrak{B}$ involves the concept of a *pseudo weak near unanimity (pwnu) polymorphism*, i.e., a polymorphism f of arity at least two such that there are endomorphisms $e_1, \dots, e_k$ of $\mathfrak{B}$ such that for all $x, y \in B$

$$e_1(f(y, x, \dots, x)) = e_2(f(x, y, \dots, x)) = \cdots = e_k(f(x, \dots, x, y)). \tag{2.1}$$

We note that all of the polynomial-time tractability conditions that we prove in this article can be phrased in terms of the existence of pwnu polymorphisms. It is not known whether every polymorphism clone of a reduct of a finitely bounded homogeneous structure that does not have a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$ contains a pwnu (see Question 21 in [19]); however, it is known to be false for polymorphism clones of ω-categorical structures in general [11]. Note that clone homomorphisms preserve identities such as (2.1), and it follows from Lemma 2.10 that primitive positive interpretations preserve the existence of pwnu polymorphisms. The same is not true for minor-preserving maps instead of clone homomorphisms. However, we have the following; it uses the assumption that $\overline{\mathrm{Aut}(\mathfrak{B})} = \mathrm{End}(\mathfrak{B})$ which is equivalent to $\mathfrak{B}$ being a *model-complete core* (see, e.g., [19, Section 4.5]).

LEMMA 2.13. *Let $\mathfrak{C}$ be a homogeneous structure with finite relational signature and let $\mathfrak{B}$ be a first-order reduct of $\mathfrak{C}$ with a pwnu polymorphism. If $\overline{\mathrm{Aut}(\mathfrak{B})} = \mathrm{End}(\mathfrak{B})$, then $\mathrm{Pol}(\mathfrak{B})$ does not have a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$.*

Proof. The assumptions imply that we may apply Corollary 3.6 in [12]. Hence, $\mathfrak{B}$ has a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$ if and only if there exist $n \in \mathbb{N}$ and $c_1, \dots, c_n \in B$ such that the clone $\mathrm{Pol}(\mathfrak{B}, \{c_1\}, \dots, \{c_n\})$ has a continuous clone homomorphism to $\mathrm{Pol}(K_3)$. But if $\mathfrak{B}$ has a pwnu polymorphism, then so does every expansion of $\mathfrak{B}$ by finitely many unary singleton relations (Proposition 10.1.13 in [19]). Since clone homomorphisms preserve the existence of pwnu polymorphisms and K_3 does not have such a polymorphism (see, e.g., [19, Proposition 6.1.43]), the statement follows. ∎

3. Algebraic Products.

We devote this section to presenting the algebraic product and studying some of its properties: Section 3.1 contains the definition together with some elementary facts while Section 3.2 describes connections with i-determined clauses (that we introduce in Section 3.1). The algebraic product has been studied in the past. For instance, Greiner [62] uses it for studying CSPs of combinations of structures or background theories (a topic we will touch upon in Section 7), Baader and Rydval [7] use it for analysing the complexity of description logics, and Bodirsky [19] uses it in connection with Ramsey structures.

3.1. Basic Properties.

The algebraic product is defined as follows.

DEFINITION 3.1. *Let $\mathfrak{A}_1$ and $\mathfrak{A}_2$ be structures with signature τ_1 and τ_2, respectively. Then the* algebraic product $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ *is the structure with domain $A_1 \times A_2$ which has for every atomic τ_1-formula $\phi(x_1, \dots, x_k)$ the relation*

$$\{((u_1, v_1), \dots, (u_k, v_k)) \mid \mathfrak{A}_1 \models \phi(u_1, \dots, u_k)\}$$

and analogously for every atomic τ_2-formula ϕ the relation

$$\{((u_1,v_1),\dots,(u_k,v_k)) \mid \mathfrak{A}_2 \models \phi(v_1,\dots,v_k)\}.$$

The relation symbol for the atomic τ_i-formula $x = y$ will be denoted by $=_i$. Clauses over the signature of $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ where all atomic formulas are built from symbols that have been introduced for atomic τ_i-formulas are called *i-determined*.

Example 3.2. The structure $(\mathbb{Q};<)\boxtimes(\mathbb{Q};<)$ has the domain $D = \mathbb{Q}^2$ and contains four relations

- $=_1$ equal to $\{((u_1,v_1),(u_2,v_2)) \in D^2 \mid u_1 = u_2\}$,
- $=_2$ equal to $\{((u_1,v_1),(u_2,v_2)) \in D^2 \mid v_1 = v_2\}$,
- $<_1$ equal to $\{((u_1,v_1),(u_2,v_2)) \in D^2 \mid u_1 < u_2\}$, and
- $<_2$ equal to $\{((u_1,v_1),(u_2,v_2)) \in D^2 \mid v_1 < v_2\}$.

The clauses $(x =_1 y \vee y <_1 z)$ and $(x <_2 y \vee x =_2 z \vee x <_2 u)$ are 1- and 2-determined, respectively. The clause $(x <_1 y \vee x =_2 y)$ is not i-determined for any $i \in \{1,2\}$.

Remark 3.3. We note that the algebraic product preserves some of the important fundamental properties of structures. For example, if $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are homogeneous, then so is $\mathfrak{A}_1\boxtimes\mathfrak{A}_2$ (Proposition 4.2.19 in [19]), and if $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are ω-categorical, then so is $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$.

We define the *n-fold algebraic product* $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$ in the natural way together with the *n-fold algebraic power*

$$\mathfrak{A}^{(n)} := \underbrace{\mathfrak{A} \boxtimes \cdots \boxtimes \mathfrak{A}}_{n \text{ times}}.$$

The forthcoming definitions and statements in this section are presented for the binary product, but they can easily be generalised to n-fold algebraic products. We continue by studying the polymorphism clone of $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. For $i \in \{1,2\}$, let $\mathscr{C}_i$ be a clone of operations on a set A_i. If $f_1 \in \mathscr{C}_1$ and $f_2 \in \mathscr{C}_2$ both have arity k, then we write (f_1,f_2) for the operation on $A_1 \times A_2$ given by

$$((a_1,b_1),\dots,(a_k,b_k)) \mapsto (f_1(a_1,\dots,a_k), f_2(b_1,\dots,b_k)).$$

The *direct product* $\mathscr{C}_1 \times \mathscr{C}_2$ of $\mathscr{C}_1$ and $\mathscr{C}_2$ is the clone $\mathscr{D}$ on the set $A_1 \times A_2$ whose operations of arity k consist of the set of all operations (f_1,f_2) where $f_i \in \mathscr{C}_i$ is of arity k. Note that this generalises the usual definition of direct products of permutation groups. If $\mathscr{C}_1 = \mathscr{C}_2 = \mathscr{C}$ then $\mathscr{D}$ is called a *direct power* and we write $\mathscr{C}^2$ instead of $\mathscr{C} \times \mathscr{C}$. Note that the function $\theta_i \colon \mathscr{C}_1 \times \mathscr{C}_2 \to \mathscr{C}_i$ given by $(f_1,f_2) \mapsto f_i$ is a uniformly continuous minor-preserving map. Also note that if $\mathscr{C}_1$ and $\mathscr{C}_2$ are oligomorphic, then so is $\mathscr{C}_1 \times \mathscr{C}_2$.

The following proposition is one of the important features of the algebraic product. We present it for two-fold algebraic products but it can obviously be generalised to the n-fold case.

Proposition 3.4. *For all structures $\mathfrak{A}_1$ and $\mathfrak{A}_2$ we have*

$$\mathrm{Pol}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2) = \mathrm{Pol}(\mathfrak{A}_1) \times \mathrm{Pol}(\mathfrak{A}_2).$$

Likewise, we have $\mathrm{End}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2) = \mathrm{End}(\mathfrak{A}_1) \times \mathrm{End}(\mathfrak{A}_1)$ *and* $\mathrm{Aut}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2) = \mathrm{Aut}(\mathfrak{A}_1) \times \mathrm{Aut}(\mathfrak{A}_2)$*.*

Proof. If $f_i \in \mathrm{Pol}(\mathfrak{A}_i)$, for $i \in \{1,2\}$, then clearly (f_1,f_2) preserves all relations of $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. Conversely, let $f \in \mathrm{Pol}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$. Pick $a \in A_2$ and define $f_1(x_1,\dots,x_n) := f((x_1,a),\dots,(x_n,a))_1$. Since f preserves $=_1$, this definition does not depend on the choice of $a \in A_2$. Note that $f_1 \in \mathrm{Pol}(\mathfrak{A}_1)$. The function $f_2 \in \mathrm{Pol}(\mathfrak{A}_2)$ is defined analogously, with the roles of 1 and 2 swapped, using a fixed element $b \in A_1$. Finally, we observe that $f = (f_1, f_2)$. Indeed, if $(x_1, y_1), \ldots, (x_n, y_n) \in A_1 \times A_2$, then

$$\begin{aligned}(f_1, f_2)((x_1, y_1), \ldots (x_n, y_n)) &= (f((x_1, a), \ldots, (x_n, a))_1, f((b, y_1), \ldots, (b, y_n))_2) \\ &= (f((x_1, y_1), \ldots, (x_n, y_n))_1, f((x_1, y_1), \ldots, (x_1, y_n))_2) \\ &= f((x_1, y_1), \ldots (x_n, y_n)),\end{aligned}$$

where the middle equality uses that f preserves $=_1$ and $=_2$. The statement for endomorphisms and automorphisms of algebraic product structures can be proved analogously. ☐

Under fairly general assumptions on $\mathscr{C}$ it holds that $\theta_i(\mathscr{C})$, for $i \in \{1, 2\}$, is closed.

Proposition 3.5. *Let $\mathfrak{A}_1$ and $\mathfrak{A}_2$ be countable ω-categorical structures and let $\mathscr{C} \subseteq \operatorname{Pol}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$ be a closed set of operations on $A_1 \times A_2$. If $\mathscr{C}$ contains αf for every $f \in \mathscr{C}$ and every $\alpha \in \operatorname{Aut}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$, then $\theta_1(\mathscr{C})$ and $\theta_2(\mathscr{C})$ are closed.*

Proof. It suffices to show the statement for $i = 1$. Let $f \in \overline{\theta_1(\mathscr{C})}$ be of arity k. Fix an enumeration $p_0, p_1, \ldots$ of A_1 and an enumeration $q_0, q_1, \ldots$ of A_2. Then for every $n \in \mathbb{N}$ there exists $g_n \in \operatorname{Pol}(\mathscr{C})$ such that $\theta_1(g_n)|_{\{p_0,\ldots,p_n\}^k} = f|_{\{p_0,\ldots,p_n\}^k}$. In the proof we use *Kőnig's tree lemma*: if T is a rooted tree with an infinite number of nodes and each node has a finite number of children, then T contains a branch of infinite length. To define the tree T, let $S_n := \{p_0, \ldots, p_n\} \times \{q_0, \ldots, q_n\}$ and consider the functions $\mathcal{F} := \bigcup_{n \in \mathbb{N}} \mathcal{F}_n$ where

$$\mathcal{F}_n := \{g_m|_{S_n^k} \mid m \geq n\}.$$

For $h_1, h_2 \in \mathcal{F}$ define $h_1 \sim h_2$ if there exists $\alpha \in \operatorname{Aut}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$ such that $h_1 = \alpha h_2$. The vertices of the tree T that we consider here are the equivalence classes of the equivalence relation $\sim$. The edges of T are defined as follows: if $h_1 \in \mathcal{F}_\ell$ is the restriction of $h_2 \in \mathcal{F}_{\ell+1}$, then the equivalence class of h_1 and the equivalence class of h_2 are linked by an edge. Clearly the tree thus defined is infinite, and by the oligomorphicity of $\operatorname{Aut}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$ there are finitely many $\sim$-classes on $\mathcal{F}_\ell$, which implies that each vertex in T has finitely many neighbours. Kőnig's tree lemma implies the existence of an infinite path in the tree.

We construct an operation g such that $g|_{S_n^k}$ is in a $\sim$-class that lies on this infinite path, for every $n \in \mathbb{N}$. This is trivial to achieve for $n = 0$. Now suppose inductively that g is already defined on S_n^k, for $n \geq 0$. By the construction of the infinite path, we find representatives h_n and h_{n+1} of the n-th and the $(n+1)$-st element on the path such that h_n is a restriction of h_{n+1}. The inductive assumption gives us an automorphism α of $\mathfrak{A}$ such that $\alpha h_n(x) = g(x)$ for all $x \in S_n^k$. We set $g(x)$ to be $\alpha h_{n+1}(x)$ for all $x \in S_{n+1}^k$. Since $\mathscr{C}$ is closed, the operation g thus defined is in $\mathscr{C}$. By the choice of g, for $n \in \mathbb{N}$, we have that $\theta_1(g) = f$. Therefore, $f \in \theta_1(\mathscr{C})$ as we wanted to prove. ☐

Another important feature of the algebraic product is that it preserves certain computational properties.

Proposition 3.6. *For $i \in \{1, 2\}$, let $\mathfrak{A}_i$ be a countable ω-categorical structure with finite relational signature τ_i. If both $\operatorname{CSP}(\mathfrak{A}_1)$ and $\operatorname{CSP}(\mathfrak{A}_2)$ are in P, then $\operatorname{CSP}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$ is in P.*

Proof. Without loss of generality, we may assume that the signatures τ_1 and τ_2 are disjoint and that $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ is a $(\tau_1 \cup \tau_2 \cup \{=_1, =_2\})$-structure. Let $\mathfrak{C}$ be an instance of $\operatorname{CSP}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$. For $i \in \{1, 2\}$, let $\mathfrak{C}_i$ be the $(\tau_i \cup \{=_i\})$-reduct of $\mathfrak{C}$. For each $i \in \{1, 2\}$, run an algorithm for $\operatorname{CSP}(\mathfrak{A}_i)$ on the input $\mathfrak{C}_i$. Accept the instance $\mathfrak{C}$ if and only if both algorithms accept.

The correctness of the algorithm follows from the fact that the map $h_i \colon C \to A_i$ is a homomorphism from $\mathfrak{C}_i$ to $\mathfrak{A}_i$ for both $i \in \{1,2\}$ if and only if the map $h \colon C \to A_1 \times A_2$ such that $h(c) = (h_1(c), h_2(c))$ for every $c \in C$ is a homomorphism from $\mathfrak{C}$ to $\mathfrak{A}_1 \times \mathfrak{A}_2$. ◻

Example 3.7. $\mathrm{CSP}(\mathbb{Q};<)$ is solvable in polynomial time by using any cycle detection algorithm for directed graphs. Hence, Proposition 3.6 implies that $\mathrm{CSP}((\mathbb{Q};<)^{(n)})$ is polynomial-time solvable for arbitrary $n \geq 1$, too.

Corollary 3.8. *Let $\mathfrak{A}_1, \mathfrak{A}_2, \mathfrak{B}$ be countable ω-categorical structures with finite relational signature such that $\mathrm{Pol}(\mathfrak{B})$ contains $\mathrm{Pol}(\mathfrak{A}_1) \times \mathrm{Pol}(\mathfrak{A}_2)$. If both $\mathrm{CSP}(\mathfrak{A}_1)$ and $\mathrm{CSP}(\mathfrak{A}_2)$ are in P, then $\mathrm{CSP}(\mathfrak{B})$ is in P, too.*

Proof. Since $\mathrm{Pol}(\mathfrak{B})$ contains $\mathrm{Pol}(\mathfrak{A}_1) \times \mathrm{Pol}(\mathfrak{A}_2) = \mathrm{Pol}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$, all relations of $\mathfrak{B}$ are primitively positively definable in $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ [37]. Hence, there is a polynomial-time reduction from $\mathrm{CSP}(\mathfrak{B})$ to $\mathrm{CSP}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$ by Lemma 2.3 so $\mathrm{CSP}(\mathfrak{B})$ is in P by Proposition 3.6. ◻

The following lemma enables us to use Lemma 2.13 for first-order expansions of algebraic products.

Lemma 3.9. *For $i \in \{1, \ldots, n\}$, let $\mathfrak{A}_i$ be a structure such that $\overline{\mathrm{Aut}(\mathfrak{A}_i)} = \mathrm{End}(\mathfrak{A}_i)$. If $\mathfrak{B}$ is a first-order expansion of $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$, then $\overline{\mathrm{Aut}(\mathfrak{B})} = \mathrm{End}(\mathfrak{B})$.*

Proof. Clearly, $\overline{\mathrm{Aut}(\mathfrak{B})} \subseteq \mathrm{End}(\mathfrak{B})$. The converse inclusion also holds since

$$\begin{aligned}
\mathrm{End}(\mathfrak{B}) \subseteq \mathrm{End}(\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n) &= \mathrm{End}(\mathfrak{A}_1) \times \cdots \times \mathrm{End}(\mathfrak{A}_n) && \text{(Proposition 3.4)} \\
&= \overline{\mathrm{Aut}(\mathfrak{A}_1)} \times \cdots \times \overline{\mathrm{Aut}(\mathfrak{A}_n)} \\
&= \overline{\mathrm{Aut}(\mathfrak{A}_1) \times \cdots \times \mathrm{Aut}(\mathfrak{A}_n)} \\
&= \overline{\mathrm{Aut}(\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n)} && \text{(Proposition 3.4)} \\
&= \overline{\mathrm{Aut}(\mathfrak{B})}.
\end{aligned}$$
◻

Lemma 3.9 holds, for instance, for the structure $(\mathbb{Q};<) \boxtimes (\mathbb{Q};<)$ since $\overline{\mathrm{Aut}(\mathbb{Q};<)}$ is equal to $\mathrm{End}(\mathbb{Q};<)$: every endomorphism of $(\mathbb{Q};<)$ is injective and preserves the complement of $<$, and by the homogeneity of $(\mathbb{Q};<)$ every restriction of an endomorphism to a finite subset of $\mathbb{Q}$ can be extended to an automorphism of $(\mathbb{Q};<)$.

We finish this section by a corollary of Lemmas 2.10, 2.13, and 3.9, which will be useful in Section 6 to prove disjointness of cases in dichotomy results.

Corollary 3.10. *Let $\mathfrak{C}$ be a first-order expansion of $(\mathbb{Q};<)^{(n)}$ and let $\mathfrak{B}$ be a structure with a pwnu polymorphism that is mutually primitively positively interpretable with $\mathfrak{C}$. Then $\mathrm{Pol}(\mathfrak{B})$ does not have a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$.*

Proof. By Lemma 2.10, there is a uniformly continuous clone homomorphism from $\mathrm{Pol}(\mathfrak{B})$ to $\mathrm{Pol}(\mathfrak{C})$. Therefore, $\mathrm{Pol}(\mathfrak{C})$ has a pwnu polymorphism. By Lemma 3.9, $\overline{\mathrm{Aut}(\mathfrak{C})} = \mathrm{End}(\mathfrak{C})$. Since $(\mathbb{Q};<)^{(n)}$ is homogeneous, $\mathrm{Pol}(\mathfrak{C})$ does not have a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$ by Lemma 2.13. Finally, by Lemma 2.10, there is a uniformly continuous clone homomorphism from $\mathrm{Pol}(\mathfrak{C})$ to $\mathrm{Pol}(\mathfrak{B})$ and therefore $\mathrm{Pol}(\mathfrak{B})$ does not have a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$ either. ◻

3.2. *i*-Determined Clauses. In the complexity analysis of first-order expansions of algebraic products $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$, we would like to use as much information about first-order expansions of $\mathfrak{A}_1$ and of $\mathfrak{A}_2$ as possible; in this context, i-determined clauses are of particular relevance. In this section we collect several general observations about definability by (conjunctions of) i-determined clauses. Throughout this section, let $\mathfrak{A}_i$ be a τ_i-structure for $i \in \{1,2\}$.

We begin by making a definition. If ϕ is a conjunction of i-determined clauses over $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$, then we let $\hat{\phi}$ denote the τ_i-formula obtained from replacing each atomic formula $R(x_1,\dots,x_k)$ by $\psi(x_1,\dots,x_k)$ where ψ is the atomic τ_i-formula for which R has been introduced in $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. Let us reconsider Example 3.2 and a formula $\phi = (x =_1 y \vee x =_1 z) \wedge (x <_1 y \vee x <_1 z \vee x <_1 u)$. Every clause in ϕ is 1-determined and $\hat{\phi}$ equals $(x = y \vee x = z) \wedge (x < y \vee x < z \vee x < u)$

LEMMA 3.11. *Let $f \in \mathrm{Pol}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$. A conjunction of i-determined clauses ϕ is preserved by f if and only if $\hat{\phi}$ is preserved by $\theta_i(f)$.*

Proof. Let $x_1,\dots,x_m$ be the free variables of ϕ. Let $((a_1^{1,1},a_2^{1,1}),\dots,(a_1^{k,m},a_2^{k,m})) \in (A_1 \times A_2)^{k\times m}$ and let f be of arity k. For $j \in \{1,\dots,k\}$, the tuple $((a_1^{j,1},a_2^{j,1}),\dots,(a_1^{j,m},a_2^{j,m}))$ satisfies ϕ in $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ if and only if $(a_i^{j,1},\dots,a_i^{j,m})$ satisfies $\hat{\phi}$ in $\mathfrak{A}_i$. Likewise,

$$\big(f((a_1^{1,1},a_2^{1,1}),\dots,(a_1^{k,1},a_2^{k,1})),\dots,f((a_1^{1,m},a_2^{1,m}),\dots,(a_1^{k,m},a_2^{k,m}))\big)$$

satisfies ϕ in $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ if and only if $\big(\theta_i(f)(a_i^{1,1},\dots,a_i^{k,1}),\dots,\theta_i(f)(a_i^{1,m},\dots,a_i^{k,m})\big)$ satisfies $\hat{\phi}$ in $\mathfrak{A}_i$. This implies the statement. ☐

We continue by analysing the polymorphisms of first-order expansion of $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ and definability via i-determined clauses. The proof is based on *reduced* formulas (see, e.g., [21, Definition 42]): a quantifier-free formula ϕ (over a structure $\mathfrak{B}$) in CNF is called reduced if all formulas obtained from ϕ by removing one of the literals from one of the clauses in the formula are not equivalent to ϕ (over $\mathfrak{B}$). Clearly, every quantifier-free formula is equivalent to a reduced formula (in particular, if a formula is unsatisfiable, the reduced formula contains an empty clause). Satisfiable reduced formulas ϕ have the property that for every literal in ϕ there exists a satisfying assignment for ϕ that satisfies the literal, but satisfies no other literal of the same clause. For example, the formula $(x < y \vee x = y) \wedge (y < x \vee z < x)$ over $(\mathbb{Q};<)$ is not reduced, but it is equivalent to the reduced formula $(x < y \vee x = y) \wedge (z < x)$.

LEMMA 3.12. *Suppose that $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are structures with quantifier elimination and let $\mathfrak{B}$ be a first-order expansion of $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. Then, the following are equivalent.*

1. *Every relation of $\mathfrak{B}$ has a definition by a conjunction of clauses each of which is either 1-determined or 2-determined.*
2. $\mathrm{Pol}(\mathfrak{B}) = \theta_1(\mathrm{Pol}(\mathfrak{B})) \times \theta_2(\mathrm{Pol}(\mathfrak{B}))$.
3. $\mathrm{Pol}(\mathfrak{B})$ *contains* (π_1^2, π_2^2).

The implication from 1 to 2 and from 2 to 3 also hold without the assumption that $\mathfrak{A}_1$ and $\mathfrak{A}_2$ have quantifier elimination.

Proof. 1 ⇒ 2. Clearly, $\mathrm{Pol}(\mathfrak{B}) \subseteq \theta_1(\mathrm{Pol}(\mathfrak{B})) \times \theta_2(\mathrm{Pol}(\mathfrak{B}))$. To prove the converse inclusion, let $g_1, g_2 \in \mathrm{Pol}(\mathfrak{B})$ be of arity k and let ϕ be a formula that defines an m-ary relation of $\mathfrak{B}$. We claim that $(\theta_1(g_1), \theta_2(g_2))$ preserves ϕ. Let $t^1 = (t_1^1,\dots,t_m^1),\dots,t^k = (t_1^k,\dots,t_m^k) \in B^m$ be tuples that satisfy ϕ in $\mathfrak{B}$. Let ψ be a conjunct of ϕ; then ψ is i-determined, for some $i \in \{1,2\}$. Since $g_i \in \mathrm{Pol}(\mathfrak{B})$ we have that $g_i(t^1,\dots,t^k)$ satisfies ψ. Since for every $j \in \{1,\dots,m\}$ we have

$$g_i(t_j^1,\dots,t_j^k)_i = (\theta_1(g_1),\theta_2(g_2))(t_j^1,\dots,t_j^k)_i$$

and ψ is i-determined, we have that $(\theta_1(g_1),\theta_2(g_2))(t^1,\dots,t^k)$ satisfies ψ as well. This implies that $(\theta_1(g_1),\theta_2(g_2))$ preserves ϕ and shows that every operation in $\theta_1(\mathrm{Pol}(\mathfrak{B})) \times \theta_2(\mathrm{Pol}(\mathfrak{B}))$ is a polymorphism of $\mathfrak{B}$.

$2 \Rightarrow 3$. Trivial.

$3 \Rightarrow 1$. We show the contrapositive. Arbitrarily choose a relation R in $\mathfrak{B}$. By assumption, R has a quantifier-free first-order definition ϕ in $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ and we may additionally assume that ϕ is written in reduced CNF. Suppose for contradiction that ϕ contains a clause which is neither 1- nor 2-determined, i.e., a clause ψ that contains a τ_1-literal ψ_1 and a τ_2-literal ψ_2. By the assumption that ϕ is reduced, ϕ has for every $i \in \{1,2\}$ a satisfying assignment α_i such that α_i satisfies ψ_i and does not satisfy all other literals of ψ. But then $x \mapsto (\pi_1^2, \pi_2^2)(\alpha_2(x), \alpha_1(x))$ satisfies none of the literals of ψ, and hence (π_1^2, π_2^2) is not in $\mathrm{Pol}(\mathfrak{B})$. ☐

The final result in this section connects primitive positive definability and i-determined clauses. We will first (in Lemmas 3.13–3.15) prove a restricted result for $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ and then extend it to n-fold products $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$ in Corollary 3.16. We use it for defining *conjunction replacement* and verify its properties; this concept is important in our algorithmic results (see Propositions 4.34 and 4.39). The proof exploits the so-called *wreath product*, which is a central group-theoretic construction. Let us denote $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ by $\mathfrak{A}$. The wreath product will be used for concisely describing the automorphism group of $\mathfrak{A}^{*j}$, where $\mathfrak{A}^{*j}$, $j \in \{1,2\}$, denotes the structure with domain $A_1 \times A_2$ that contains all relations that are defined by a j-determined clause. If G is a permutation group on A_1 and H is a permutation group on A_2, then the *action of the (unrestricted) wreath product $G \ltimes H^{A_1}$ on $A_1 \times A_2$* (also denoted in the literature by $H \operatorname{Wr}_{A_1} G$) is the permutation group

$$\{(a_1, a_2) \mapsto (\alpha(a_1), \beta_{a_1}(a_2)) \mid \alpha \in G, \beta_a \in H \text{ for every } a \in A_1\}.$$

Lemma 3.13. *For any structures $\mathfrak{A}_1$ and $\mathfrak{A}_2$, let $\mathfrak{A}$ denote $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. Then it holds that*

$$\mathrm{Aut}(\mathfrak{A}^{*1}) = \mathrm{Aut}(\mathfrak{A}_1) \ltimes \mathrm{Sym}(A_2)^{A_1}$$

and if $\mathfrak{A}_1$ is homogeneous then so is $\mathfrak{A}^{*1}$. The analogous statements hold for $\mathfrak{A}^{*2}$.

Proof. To show that $\mathrm{Aut}(\mathfrak{A}_1) \ltimes \mathrm{Sym}(A_2)^{A_1} \subseteq \mathrm{Aut}(\mathfrak{A}^{*1})$, let $\psi(x_1, \ldots, x_m)$ be a 1-determined clause and let $((s_1,t_1), \ldots, (s_m,t_m))$ be a tuple that satisfies ψ; i.e., there exists an atomic τ_1-formula ϕ such that $(s_1, \ldots, s_m)$ satisfies ϕ. For $\alpha \in \mathrm{Aut}(\mathfrak{A}_1)$ and $\beta_{s_1}, \ldots, \beta_{s_m} \in \mathrm{Sym}(A_2)$, note that

$$((\alpha(s_1), \beta_{s_1}(t_1)), \ldots, (\alpha(s_m), \beta_{s_m}(t_m)))$$

satisfies ψ since $(\alpha(s_1), \ldots, \alpha(s_m))$ satisfies ϕ.

To show that $\mathrm{Aut}(\mathfrak{A}^{*1}) \subseteq \mathrm{Aut}(\mathfrak{A}_1) \ltimes \mathrm{Sym}(A_2)^{A_1}$, let $\gamma \in \mathrm{Aut}(\mathfrak{A}^{*1})$. Arbitrarily fix $t \in A_2$. The operation γ preserves $=_1$ so the operation α defined by $s \mapsto \gamma((s,t))_1$ is well-defined and it is an automorphism of $\mathfrak{A}_1$. The operation γ is bijective so for every $s \in A_1$, the map β_s defined by $t \mapsto \gamma(s,t)_2$ is a member of $\mathrm{Sym}(A_2)$. Since γ equals the map that sends (s,t) to $(\alpha(s), \beta_s(t))$, this shows that $\gamma \in \mathrm{Aut}(\mathfrak{A}_1) \ltimes \mathrm{Sym}(A_2)^{A_1}$.

Now suppose that $\mathfrak{A}_1$ is homogeneous. Let α be an isomorphism between m-element substructures of $\mathfrak{A}_1^{*1}$ that maps (s_j, t_j) to (s'_j, t'_j) for $j \in \{1, \ldots, m\}$ and $m \in \mathbb{N}$. Note that if $s_j = s_k$, then $s'_j = s'_k$ because α must preserve the relation $=_1$. Hence, the map α_1 that sends s_j to s'_j is a well-defined map between finite subsets of A_1. Moreover, since α preserves all i-determined clauses, the map α_1 is in fact an isomorphism between finite substructures of $\mathfrak{A}_1$, and hence can be extended to an automorphism β of $\mathfrak{A}_1$ by the homogeneity of $\mathfrak{A}_1$. Note that if p,q are distinct and $s_p \neq s_q$, then $\alpha(s_p,t_p) \neq \alpha(s_q,t_q)$, because α is injective. Hence, for each s in the domain of α_1 we may fix a bijection γ_s of A_2 such that $\alpha(s,t) = (\beta(s), \gamma_s(t))$ for all t such that (s,t) lies in the domain of α. For all other $s \in A_1$ we may define γ_s to be the identity. Then the map that sends (a,b) to $(\beta(a), \gamma_a(b)) \in \mathrm{Aut}(\mathfrak{A}^{*1})$ extends α. We conclude that $\mathfrak{A}^{*1}$ is homogeneous. ☐

Lemma 3.14. *Assume that $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are countable, homogeneous, and ω-categorical. A relation R can be defined by a conjunction of 1-determined clauses over $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ if and only if it is preserved by the wreath product*

$$\mathrm{Aut}(\mathfrak{A}_1) \ltimes \mathrm{Sym}(A_2)^{A_1}$$

in its action on $A_1 \times A_2$; the analogous characterisation holds for clauses that are 2-determined.

Proof. The forward implication is an immediate consequence of Lemma 3.13. Conversely, suppose that R is preserved by $\mathrm{Aut}(\mathfrak{A}_1) \ltimes \mathrm{Sym}(A_2)^{A_1}$. Recall that $\mathrm{Aut}(\mathfrak{A}^{*1}) = \mathrm{Aut}(\mathfrak{A}_1) \ltimes \mathrm{Sym}(A_2)^{A_1}$ by Lemma 3.13. The structure $\mathfrak{A}_1$ is homogeneous by assumption so $\mathfrak{A}^{*1}$ is homogeneous, too. We have assumed that $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are ω-categorical so $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ is ω-categorical by Remark 3.3. Consequently, $\mathfrak{A}^{*1}$ is ω-categorical since it is a first-order reduct of $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. It follows that R is first-order definable in $\mathfrak{A}^{*1}$, and even has a quantifier-free definition because $\mathfrak{A}^{*1}$ is homogeneous. This implies that R can be defined by a conjunction of 1-determined clauses over $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. ◻

Lemma 3.15. *Assume the following:*

1. *$\mathfrak{A}_1$ and $\mathfrak{A}_2$ are countable, homogeneous, and ω-categorical,*
2. *$\mathfrak{B}$ is a first-order expansion of $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$, and*
3. *$\phi_1 \wedge \phi_2$ is a formula that defines a relation R over $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$ such that ϕ_1 is a conjunction of 1-determined clauses and such that R is primitively positively definable over $\mathfrak{B}$.*

Then there exists a conjunction ψ_1 of 1-determined clauses which is equivalent to a primitive positive formula over $\mathfrak{B}$ such that $\psi_1 \wedge \phi_2$ still defines R over $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. An analogous statement holds when ϕ_2 is a conjunction of 2-determined clauses.

Proof. Suppose that items 1–3 hold true. Let $\psi(x_1, \dots, x_m)$ be the formula

$$\exists y_1, \dots, y_m \left(R(y_1, \dots, y_m) \wedge \bigwedge_{j \in \{1,\dots,m\}} x_j =_1 y_j \right).$$

Note that ψ is equivalent to a primitive positive formula over $\mathfrak{B}$.

We first show that the relation S defined by ψ over $\mathfrak{B}$ can be defined by a conjunction of i-determined clauses over $\mathfrak{A}_1 \boxtimes \mathfrak{A}_2$. We use Lemma 3.14. Let $((a_1^1, a_2^1), \dots, (a_1^m, a_2^m)) \in S$, let $\alpha \in \mathrm{Aut}(\mathfrak{A}_1)$, and let $\pi_1, \dots, \pi_m \in \mathrm{Sym}(A_2)$ be such that $\pi_p = \pi_q$ whenever $a_1^p = a_1^q$ for $p, q \in \{1, \dots, m\}$. We have to show that

$$t := \big((\alpha(a_1^1), \pi_1(a_2^1)), \dots, (\alpha(a_1^m), \pi_m(a_2^m))\big)$$

satisfies ψ as well. Let (b_1^1, b_2^1), …, $(b_1^m, b_2^m) \in B$ be the witnesses from B for the existentially quantified variables of ψ that show that ψ holds for $((a_1^1, a_2^1), \dots, (a_1^m, a_2^m))$. Then the tuple $((\alpha(b_1^1), b_2^1), \dots, (\alpha(b_1^m), b_2^m))$ provides witnesses that show that the same formula holds for t:

- $R((\alpha(b_1^1), b_2^1), \dots, (\alpha(b_1^m), b_2^m))$ holds in $\mathfrak{B}$ because $R((b_1^1, b_2^1), \dots, (b_1^m, b_2^m))$ holds in $\mathfrak{B}$ and $(\alpha, \mathrm{id}) \in \mathrm{Aut}(\mathfrak{B})$, and
- $(\alpha(a_1^j), \pi(a_2^j)) =_1 (\alpha(b_1^j), b_2^j)$ holds for every $j \in \{1, \dots, m\}$ because $(a_1^j, a_2^j) =_1 (b_1^j, b_2^j)$.

Lemma 3.14 shows that there exists a conjunction ψ_1 of 1-determined clauses that is equivalent to ψ. Clearly, ψ_1 implies ϕ_1 (since ϕ_1 is 1-determined), so $\psi_1 \wedge \phi_2$ defines R in $\mathfrak{B}$ and this concludes the proof.

The case when ϕ_2 is a conjunction of 2-determined clauses can be treated analogously. ◻

We continue by generalising the previous lemma to algebraic products involving more than two structures. To this end, we need a particular notion that generalizes i-determined clauses. Let

$S \subseteq \{1, \ldots, n\}$. A clause over $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$ is called *S-determined* if all atomic formulas in the clause are built from symbols that have been introduced for atomic τ_i-formulas for some $i \in S$.

Corollary 3.16. *Assume the following:*

1. *$\mathfrak{A}_1, \ldots, \mathfrak{A}_n$ are countable, homogeneous, and ω-categorical,*
2. *$\mathfrak{B}$ is a first-order expansion of $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$, and*
3. *$\phi_S \wedge \phi$ is a formula that defines a relation R over $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$ such that for some $S \subseteq \{1, \ldots, n\}$, the formula ϕ_S is a conjunction of S-determined clauses and such that R is primitively positively definable over $\mathfrak{B}$.*

Then there exists a conjunction ψ_S of S-determined clauses which is equivalent to a primitive positive formula over $\mathfrak{B}$ such that $\psi_S \wedge \phi$ still defines R over $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$.

Proof. Assume without loss of generality that $S = \{1, \ldots, p\}$ for some $p \geq 1$. We can view the n-fold product $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$ as $\mathfrak{B}_1 \boxtimes \mathfrak{B}_2$, where $\mathfrak{B}_1 = \mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_p$ and $\mathfrak{B}_2 = \mathfrak{A}_{p+1} \boxtimes \cdots \boxtimes \mathfrak{A}_n$. Note that ϕ_1 is a conjunction of 1-determined clauses when considered as a formula over $\mathfrak{B}_1 \boxtimes \mathfrak{B}_2$. By Lemma 3.15, there exists a conjunction ψ_S of 1-determined clauses which is equivalent to a primitive positive formula over $\mathfrak{B}$ such that $\psi_S \wedge \phi$ still defines R. Since ψ_S is S-determined when viewed as a formula over $\mathfrak{A}_1 \boxtimes \cdots \boxtimes \mathfrak{A}_n$, the claim follows. ∎

Let $\mathfrak{A}_1, \ldots \mathfrak{A}_n$, $\mathfrak{B}$, S, and $\phi_S \wedge \phi$ be as in the statement of Corollary 3.16. Arbitrarily choose a conjunction ψ_S of S-determined clauses equivalent to a primitive positive formula over $\mathfrak{B}$ such that $\psi_S \wedge \phi$ is equivalent to $\phi_S \wedge \phi$. Note that the existence of ψ_S follows from the corollary. We denote the formula ψ_S by $\mathrm{cr}(\phi_S \wedge \phi, S, \phi_S)$, where cr stands for *conjunction replacement*. Note that since ψ_S is equivalent to a primitive positive formula, it is preserved by $\mathrm{Pol}(\mathfrak{B})$, which will be relevant for proving syntactic restrictions on ψ_S and use of known algorithms for temporal CSPs.

Remark 3.17. In a typical example of a formula $\phi_S \wedge \phi$, where ϕ_S is a conjunction of all S-determined clauses of $\phi_S \wedge \phi$, conjunction replacement is not needed, because ϕ_S is already preserved by $\mathrm{Pol}(\mathfrak{B})$ and therefore equivalent to a primitive positive formula over $\mathfrak{B}$. This is because syntactic normal forms are often defined as conjunctions of clauses of some specific shape. Nevertheless, one cannot guarantee that a general defining first-order formula of R will be of this shape and therefore ϕ_S might not be preserved by $\mathrm{Pol}(\mathfrak{B})$. Since we want to exploit the syntactic normal forms proved for temporal CSPs (see Section 4.1), we ensure this property using the operator cr.

4. Algebraic Powers of $(\mathbb{Q}; <)$.

In this section we classify the complexity of the CSP for every first-order expansion of the structure $(\mathbb{Q}; <)^{(n)}$. The main outline of the section is as follows: we consider the case when $n = 2$ in Section 4.4, generalize to arbitrary $n \geq 2$ in Section 4.5, and finally specialize our results to binary signatures in Section 4.6. We begin the section by recapitulating some known results concerning first-order expansions of $(\mathbb{Q}; <)$ (Section 4.1), studying the polymorphisms of $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$ (Section 4.2) and presenting syntactic normal forms of certain relations that are first-order definable in $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$ (Section 4.3).

Recall that $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <) = (\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ is ω-categorical and homogeneous (Remark 3.3), and therefore has quantifier elimination. From here until Section 4.4, we let the symbol $\mathfrak{D}$ denote a first-order expansion of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$.

4.1. First-order Expansions of $(\mathbb{Q}; <)$.

Let $\mathfrak{B}$ be a first-order reduct of $(\mathbb{Q}; <)$ with a finite relational signature. The complexity of $\mathrm{CSP}(\mathfrak{B})$ for all choices of $\mathfrak{B}$ has been determined by Bodirsky and Kára [30]. For our purposes, it is sufficient to understand the complexity of all first-order expansions of $(\mathbb{Q}; <)$ with a finite relational structure. We next present first-order expansions of $(\mathbb{Q}; <)$ with a polynomial-time solvable CSP, we describe some of their polymorphisms, and how the relations can be described with syntactically restricted definitions. We make use of relational and functional dualities to simplify the presentation.

Definition 4.1. *The* dual *of a relation* $R \subseteq \mathbb{Q}^k$ *is the relation*

$$R^* := \{(-x_1, \dots, -x_k) \mid (x_1, \dots, x_k) \in R\}.$$

If $\mathfrak{B}$ *is a relational structure with domain* $\mathbb{Q}$*, then the* dual *of* $\mathfrak{B}$ *is the structure with domain* $\mathbb{Q}$ *and the same signature* τ *as* $\mathfrak{B}$ *where* $R \in \tau$ *denotes* $(R^{\mathfrak{B}})^*$*. Similarly, if* $f\colon \mathbb{Q}^n \to \mathbb{Q}$ *is an operation, then the* dual *of* f *is the operation* f^* *defined as follows.*

$$(x_1, \dots, x_n) \mapsto -f(-x_1, \dots, -x_n)$$

If $\mathscr{C}$ *is a operation clone on* $\mathbb{Q}$*, then the* dual *of* $\mathscr{C}$ *is the operation clone* $\mathscr{C}^* := \{f^* \mid f \in \mathscr{C}\}$*.*

Now, consider the following first-order expansions of $(\mathbb{Q}; <)$.

- $\mathfrak{U} := (\mathbb{Q}; <, R^{\min}_{\leq})$ where $R^{\min}_{\leq} := \{(x, y, z) \in \mathbb{Q}^3 \mid y \leq x \text{ or } z \leq x\}$.
- $\mathfrak{X} := (\mathbb{Q}; <, X)$ where $X := \{(x, y, z) \in \mathbb{Q}^3 \mid x = y < z \text{ or } y = z < x \text{ or } z = x < y\}$
- $\mathfrak{I} := (\mathbb{Q}; <, R^{\text{mi}}, S^{\text{mi}})$ where

$$R^{\text{mi}} := \{(x, y, z) \in \mathbb{Q}^3 \mid y \leq x \text{ or } z < x\} \text{ and}$$
$$S^{\text{mi}} := \{(x, y, z) \in \mathbb{Q}^3 \mid y \neq x \text{ or } z \leq x\}.$$

- $\mathfrak{L} := (\mathbb{Q}; <, L, I_4)$ where

$$L := \{(x, y, z) \in \mathbb{Q}^3 \mid y < x \text{ or } z < x \text{ or } x = y = z\} \text{ and}$$
$$I_4 := \{(x, y, u, v) \in \mathbb{Q}^4 \mid x = y \text{ implies } u = v\}.$$

We need the following characterisation of primitive positive definability in these structures in terms of certain polymorphisms. The precise definition of these operations can be found in [39] but it is not needed in this article; the properties stated in the next proposition suffice for our purposes.

Theorem 4.2 (Proposition 7.27 in [39]). *For every* $k \geq 3$ *there are pseudo weak near unanimity polymorphisms* $\text{min}_k, \text{mx}_k, \text{mi}_k, \text{ll}_k$ *of arity* k *such that a relation* $R \subseteq \mathbb{Q}^m$ *is preserved by* min_k *(*mx_k*,* mi_k*,* ll_k*) and* $\text{Aut}(\mathbb{Q}; <)$ *if and only if* R *has a primitive positive definition in* $\mathfrak{U}$ *(*$\mathfrak{X}$*,* $\mathfrak{I}$*,* $\mathfrak{L}$*). The operation* ll_k *is injective.*

The following result is essentially taken from [30] but we formulate it differently with the aid of polymorphisms.

Theorem 4.3 (Theorem 12.0.1 in [19]). *Let* $\mathfrak{B}$ *be a first-order expansion of* $(\mathbb{Q}; <)$*. Then exactly one of the following two cases applies.*

1. $\mathfrak{B}$ *is preserved by the operation* min_3*,* mx_3*,* mi_3*, or* ll_3 *from Theorem 4.2, or the dual of one of these operations. In this case, the CSP of every finite-signature reduct of* $\mathfrak{B}$ *is in P.*
2. $\text{Pol}(\mathfrak{B})$ *has a uniformly continuous minor-preserving map to* $\text{Pol}(K_3)$*. In this case,* $\mathfrak{B}$ *has a finite-signature reduct whose CSP is NP-complete.*

Theorem 4.3 immediately connects first-order expansions of $(\mathbb{Q}; <)$ with the infinite-domain tractability conjecture from Section 2.3.

We now describe polymorphisms and syntactic normal forms of the structures that were described earlier. Clearly, an operation f preserves a relation $R \subseteq \mathbb{Q}^k$ if and only if f^* preserves R^* so we can concentrate on the structures $\mathfrak{U}$, $\mathfrak{X}$, $\mathfrak{I}$, and $\mathfrak{L}$. We start by considering polymorphisms of the structures $\mathfrak{U}$, $\mathfrak{X}$, and $\mathfrak{I}$.

DEFINITION 4.4. *A binary operation f on $\mathbb{Q}$ is called a* pp-operation *if $f(a,b) \leq f(a',b')$ if and only if*

1. $a \leq 0$ *and* $a \leq a'$, *or*
2. $0 < a$, $0 < a'$, *and* $b \leq b'$.

Remark 4.5. Note that if $f, g \colon \mathbb{Q}^k \to \mathbb{Q}$ are such that for all $a, b \in \mathbb{Q}^k$ we have $f(a) \leq f(b) \Leftrightarrow g(a) \leq g(b)$, then a relation R which is first-order definable in $(\mathbb{Q}; <)$ is preserved by f if and only if it is preserved by g. To see this, suppose that R is of arity k and $a^1, \dots, a^k \in R$. Then $s \coloneqq f(a^1, \dots, a^k)$ and $t \coloneqq g(a^1, \dots, a^k)$ satisfy the same atomic formulas over $(\mathbb{Q}; <)$, and hence by the homogeneity of $(\mathbb{Q}; <)$ there exists $\alpha \in \mathrm{Aut}(\mathbb{Q}; <)$ which maps s to t, and since R is first-order definable over $(\mathbb{Q}; <)$ either both or neither of s and t lie in R. In particular, R is preserved by a pp-operation if and only if it is preserved by all pp-operations.

PROPOSITION 4.6 ([30]). *Each of the structures $\mathfrak{U}$, $\mathfrak{X}$, and $\mathfrak{I}$ is preserved by a* pp*-operation. Equivalently, if a relation $R \subseteq \mathbb{Q}^k$ with a first-order definition in $(\mathbb{Q}; <)$ is preserved by* min_3, mx_3 *or* mi_3, *then it is preserved by a* pp*-operation.*

One should note that if a structure $\mathfrak{B}$ is preserved by a pp-operation, then this does not imply that CSP($\mathfrak{B}$) is polynomial-time solvable. It does, however, imply that the relations in $\mathfrak{B}$ can be defined via a restricted form of definitions.

THEOREM 4.7 (Theorem 4 in [22]). *Let $R \subseteq \mathbb{Q}^m$ be a relation with a first-order definition in $(\mathbb{Q}; <)$. Then the following are equivalent.*

- *R is preserved by a (equivalently: every)* pp*-operation.*
- *R has a definition by a conjunction of clauses of the form*

$$y_1 \neq x \vee \cdots \vee y_k \neq x \vee z_1 \leq x \vee \cdots \vee z_l \leq x$$

  where it is permitted that $l = 0$ or $k = 0$.

We conclude this section by another characterisation of $\mathfrak{L}$ via polymorphisms and presenting a syntactic normal form.

DEFINITION 4.8. *A binary operation f on $\mathbb{Q}$ is called an* ll-operation *if $f(a,b) < f(a',b')$ if and only if*

1. $a \leq 0$ *and* $a < a'$, *or*
2. $a \leq 0$ *and* $a = a'$ *and* $b < b'$, *or*
3. $a, a' > 0$ *and* $b < b'$, *or*
4. $a > 0$ *and* $b = b'$ *and* $a < a'$.

Note that every ll-operation is injective; this fact will be important in some of the forthcoming proofs. A visualisation of a pp-operation and an ll-operation can be found in Figure 4.1.

DEFINITION 4.9. *A formula is an* ll-Horn clause *if it is of the form*

$$x_1 \neq y_1 \vee \cdots \vee x_m \neq y_m \vee z_1 < z_0 \vee \cdots \vee z_\ell < z_0$$

or

$$x_1 \neq y_1 \vee \cdots \vee x_m \neq y_m \vee z_1 < z_0 \vee \cdots \vee z_\ell < z_0 \vee (z_0 = z_1 = \cdots = z_\ell)$$

where it is permitted that $l = 0$ or $m = 0$.

We also need *lexicographic operations* in order to formulate the final theorem of this section.

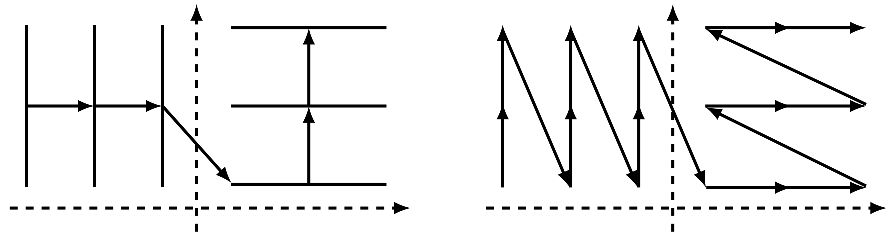

FIG. 4.1. *A visualisation of a pp-operation (left) and an* ll*-operation (right) [29]. Arrows depict the growth of values.*

TABLE 4.1
Summary of Definitions 4.4, 4.8, and 4.10.

f is pp if	$f(a,b) \leq f(a',b')$	$\iff$	$(a \leq 0 \wedge a \leq a') \vee (0 < a \wedge 0 < a' \wedge b \leq b')$
f is ll if	$f(a,b) < f(a',b')$	$\iff$	$(a \leq 0 \wedge a < a') \vee (a \leq 0 \wedge a = a' \wedge b < b') \vee$ $(a > 0 \wedge a' > 0 \wedge b < b') \vee (a > 0 \wedge b = b' \wedge a < a')$
f is lex if	$f(a,b) < f(a',b')$	$\iff$	$(a < a') \vee (a = a' \wedge b < b')$

DEFINITION 4.10. *A binary operation f on $\mathbb{Q}$ is called a* lex-operation *if $f(a,b) < f(a',b')$ if and only if*
- *$a < a'$, or*
- *$a = a'$ and $b < b'$.*

It is called a twisted lex-operation *if $f(a,-b)$ is a* lex*-operation.*

Remark 4.11. Every relation $R \subseteq \mathbb{Q}^k$ with a first-order definition in $(\mathbb{Q};<)$ that is preserved by an ll-operation is also preserved by all lex-operations.

THEOREM 4.12 ([32] and [81]; also see Theorem 12.7.3 and Lemma 12.4.4 in [19]). *Let $R \subseteq \mathbb{Q}^k$ be a relation with a first-order definition in $(\mathbb{Q};<)$. Then the following are equivalent.*
- *R has a primitive positive definition in $\mathfrak{L}$.*
- *R is preserved by an (equivalently: every)* ll*-operation.*
- *R is preserved by ll_k (from Theorem 4.2) for some (equivalently: for all) $k \geq 3$.*
- *R has a definition by a conjunction of ll-Horn clauses.*

Moreover, if R is preserved by a pp*-operation and by a* lex*-operation, then R is preserved by an* ll*-operation.*

The operations discussed in this section will be used frequently in the sequel. A concise summary can be found in Table 4.1.

4.2. Polymorphisms. We will now analyse the polymorphism clones of first-order expansions of $(\mathbb{Q};<)\boxtimes(\mathbb{Q};<)$. This involves a study of *canonical functions* (see, e.g., [43]) in the product setting.

We formulate the results for the polymorphisms of first-order expansions of $(\mathbb{Q};<) \boxtimes (\mathbb{Q};<)$, since they can be described more explicitly than in first-order expansions of $(\mathbb{Q};<)^{(n)}$. Nevertheless, generalised versions of Lemma 4.15, Lemma 4.17, Lemma 4.18, Corollary 4.19 and Lemma 4.20 for first-order expansions of $(\mathbb{Q};<)^{(n)}$ can be proved in a similar fashion. The basic idea is to choose one or two dimensions from $\{1,\dots,n\}$ that are referred to in the statements. For a concrete example, see the proof of Proposition 4.37 that is a generalization of Corollary 4.19. Generalizations of this kind will be important in Section 4.5.

Let G be a permutation group on a set A and let H be a permutation group on a set B. A function $f\colon A \to B$ is called *canonical with respect to* (G,H) if for every $m \in \mathbb{N}$, $t \in A^m$, and $\alpha \in G$ there exists a $\beta \in H$ such that $f\alpha(t) = \beta f(t)$ (where functions are applied to tuples componentwise). If f is canonical with respect to $(\mathrm{Aut}(\mathfrak{A})^n, \mathrm{Aut}(\mathfrak{A}))$ for some $n \in \mathbb{N}$, then we say that f is *canonical over* $\mathrm{Aut}(\mathfrak{A})$. In other words, f is canonical over $\mathrm{Aut}(\mathfrak{A})$ if and only if for every $m \in \mathbb{N}$ and all $t_1,\dots,t_n \in A^m$ the orbit of $f(t_1,\dots,t_n)$ in $\mathrm{Aut}(\mathfrak{A})$ only depends on the orbits of $t_1,\dots,t_n$ in $\mathrm{Aut}(\mathfrak{A})$. Note that if $\mathfrak{B} = \mathfrak{A}_1 \boxtimes \mathfrak{A}_2$, then an operation f is canonical over $\mathrm{Aut}(\mathfrak{B})$ if $\theta_1(f)$ is canonical over $\mathfrak{A}_1$ and $\theta_2(f)$ is canonical over $\mathfrak{A}_2$.

Example 4.13. Let f be a lex-operation and g a twisted lex-operation (see Definition 4.10). Then f and g are canonical over $\mathrm{Aut}(\mathbb{Q};<)$ and (f,g) is canonical over $\mathrm{Aut}((\mathbb{Q};<) \boxtimes (\mathbb{Q};<))$.

A permutation group G is called *extremely amenable* if every continuous action of G on a compact Hausdorff space has a fixed point. The reader need not be familiar with this notion since it will only be used in a black-box fashion via Theorem 4.14 below; we refer the interested reader to [68]. A fundamental example of a structure with an extremely amenable automorphism group is $(\mathbb{Q};<)$. Moreover, direct products of extremely amenable groups are extremely amenable [68].

Theorem 4.14 (see, e.g., [45, 43]). *Let G be an extremely amenable permutation group on a set A, let H be an oligomorphic permutation group on a set B, and let $f\colon A \to B$ be a function. Then*

$$\overline{\{\beta f \alpha \mid \alpha \in G, \beta \in H\}}$$

contains a canonical function with respect to (G,H).

The following result will be useful later on when we analyse the polymorphisms of first-order expansions of powers of $(\mathbb{Q};<)$. If f is an operation of arity k and $\alpha_1,\dots,\alpha_n$ are unary operations, then we write $f(\alpha_1,\dots,\alpha_n)$ to denote the function $(x_1,\dots,x_n) \mapsto f(\alpha_1(x_1),\dots,\alpha_n(x_n))$.

Lemma 4.15. *Let $\mathfrak{A}_1, \mathfrak{A}_2$ be ω-categorical structures such that $\mathrm{Aut}(\mathfrak{A}_1)$ is extremely amenable and assume $f \in \mathrm{Pol}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$ has arity n. Then, the set*

$$\mathscr{C} := \overline{\{\alpha_0 f(\alpha_1,\dots,\alpha_n) \mid \alpha_0 \in \mathrm{Aut}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2),\ \alpha_j \in \mathrm{Aut}(\mathfrak{A}_1) \times \{\mathrm{id}_{A_2}\} \text{ for all } j = 1,2,\dots,n\}}$$

contains an operation g such that $\theta_1(g)$ is canonical over $\mathrm{Aut}(\mathfrak{A}_1)$, and $\theta_2(g) = \theta_2(f)$. The symmetric statement holds if the roles of $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are exchanged.

Proof. By Theorem 4.14 there exists an operation

$$g'' \in \overline{\{\alpha_0\theta_1(f)(\alpha_1,\dots,\alpha_n) \mid \alpha_0,\alpha_1,\dots,\alpha_n \in \mathrm{Aut}(\mathfrak{A}_1)\}}$$

which is canonical over $\mathrm{Aut}(\mathfrak{A}_1)$. Note that $g'' \in \overline{\theta_1(\mathscr{C})}$. By Proposition 3.5, $\theta_1(\mathscr{C})$ is closed and therefore there exists $g' \in \mathscr{C}$ such that $\theta_1(g') = g''$. Arbitrarily choose $a^1,\dots,a^k$ in A_2^n. The definition of $\mathscr{C}$ implies that there is an automorphism $\alpha \in \mathrm{Aut}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)$ such that $\theta_2(g')(a^i) = \theta_2(\alpha)\theta_2(f)(a^i)$ for $i = 1,\dots,k$, which we can rewrite as $\theta_2(\alpha^{-1})\theta_2(g')(a^i) = \theta_2(f)(a^i)$. This shows that $\theta_2(f) \in \overline{\{\beta\theta_2(g') \mid \beta \in \mathrm{Aut}(\mathfrak{A}_2)\}}$. Let $S := \overline{\{\gamma g' \mid \gamma \in \mathrm{Aut}(\mathfrak{A}_1 \boxtimes \mathfrak{A}_2)\}}$. Note that $\theta_2(S)$ contains $\{\beta\theta_2(g') \mid \beta \in \mathrm{Aut}(\mathfrak{A}_2)\}$ so $\overline{\theta_2(S)}$ contains $\theta_2(f)$. Applying Proposition 3.5 to the set S implies that $\theta_2(S)$ is closed. Hence, there exists $g \in S$ such that $\theta_2(g) = \theta_2(f)$. Note that for every finite subset F of A_1, $\theta_1(g)|_F = \gamma' g''|_F$ for some $\gamma' \in \mathrm{Aut}(\mathfrak{A}_1)$. Therefore, $\theta_1(g)$ is canonical over $\mathrm{Aut}(\mathfrak{A}_1)$, because g'' is. ∎

We continue to introduce terminology. An operation $f\colon A^k \to A$ is called *essentially unary* if there exists $i \in \{1,\dots,k\}$ and a unary operation $g\colon A \to A$ such that $f(x_1,\dots,x_k) = g(x_i)$ for all $x_1,\dots,x_k \in A$. Let $S \subseteq \mathbb{Q}$. An operation $f\colon \mathbb{Q}^2 \to \mathbb{Q}$ is called *dominated by the first argument on* S if $f(x,y) < f(x',y')$ for all $x,x' \in S$ such that $x < x'$. If an operation $f\colon \mathbb{Q}^2 \to \mathbb{Q}$ is dominated by the first argument on all of $\mathbb{Q}$, we say that is *dominated by the first argument*. Examples of operations that are dominated by their first argument are lex-operations, twisted lex-operations, and order-preserving operations that only depend on the first argument.

Recall that $\mathfrak{D}$ is an arbitrary fixed first-order expansion of $(\mathbb{Q};<)\boxtimes(\mathbb{Q},<) = (\mathbb{Q};<_1,=_1,<_2,=_2)$. Our aim is now to show that $\mathrm{Pol}(\mathfrak{D})$ contains an operation with suitable domination properties (Lemma 4.18). This lemma will be a cornerstone in the proof of our first result on syntactic normal forms (Proposition 4.26). We first note that the binary polymorphisms of $(\mathbb{Q};<)$ that are canonical over $\mathrm{Aut}(\mathbb{Q};<)$ can be given a succinct characterisation.

Lemma 4.16 (see, e.g., Example 11.4.13 in [19]). *Assume that $f \in \mathrm{Pol}(\mathbb{Q};<)$ is a binary operation that is canonical over $\mathrm{Aut}(\mathbb{Q};<)$. Then, either f is essentially unary, or f is a* lex-*operation or a twisted* lex-*operation, or the operation $(x,y) \mapsto f(y,x)$ is a* lex-*operation or a twisted* lex-*operation.*

We now turn our attention to the structure $\mathfrak{D}$ and obtain the following intermediate result by analysing operations g in $\mathrm{Pol}(\mathfrak{D})$ that are canonical in a particular dimension. Note that the following statements of Lemma 4.17, Lemma 4.18 and Corollary 4.19 remain true also if the duals of ll- or pp-operations are used.

Lemma 4.17. *If $\mathrm{Pol}(\mathfrak{D})$ contains an operation f such that $\theta_1(f)$ is an* ll-*operation, then $\mathrm{Pol}(\mathfrak{D})$ also contains an operation g such that $\theta_1(g)$ is an* ll-*operation and $\theta_2(g)$ or $(x,y) \mapsto \theta_2(g)(y,x)$ is either a* lex-*operation or essentially unary (and in particular preserves $\leq_2$ and $\neq_2$). The analogous statement holds if $\theta_1(f)$ is a* pp-*operation.*

Proof. Apply Lemma 4.15 to the operation f for dimension $i = 2$ and let $g \in \mathrm{Pol}(\mathfrak{D})$ be the resulting operation such that $\theta_2(g)$ is canonical and $\theta_1(g) = \theta_1(f)$. By Lemma 4.16, either $\theta_2(g)$ is essentially unary, or a lex-operation, or a twisted lex-operation, or $(x,y) \mapsto \theta_2(g)(y,x)$ is a lex-operation, or a twisted lex-operation. If $\theta_2(g)$ is a twisted lex-operation, then we consider g' defined by $g'(x,y) \coloneqq g(x,g(x,y))$ which is a lex-operation. The argument if $(x,y) \mapsto \theta_2(g)(y,x)$ is a twisted lex-operation is similar. We finally note that $\theta_1(g')$ is an ll-operation. The same proof works if $\theta_1(f)$ is a pp-operation. ∎

In our final step, we show that if $\mathrm{Pol}(\mathfrak{D})$ contains an operation that satisfies the preconditions of Lemma 4.17, then the expanded structure $(\mathfrak{D};\leq_1,\neq_1,\leq_2,\neq_2)$ admits a polymorphism with a certain domination property.

Lemma 4.18. *Let $f \in \mathrm{Pol}(\mathfrak{D})$ be such that $\theta_1(f)$ is an* ll-*operation or a* pp-*operation and let $i \in \{1,2\}$. Then $\mathrm{Pol}(\mathfrak{D};\leq_1,\neq_1)$ contains an operation g such that $\theta_2(g) = \theta_2(f)$ and $\theta_1(g)$ is dominated by the i-th argument. If $\theta_1(f)$ is a* pp-*operation, then g can be chosen such that $\theta_1(g)$ equals π^2_i.*

Proof. We begin with the case $i = 1$. Define

$$U = \overline{\{\beta f(\alpha, \mathrm{id}_{\mathbb{Q}^2}) \mid \alpha,\beta \in \mathrm{Aut}(\mathbb{Q};<) \times \{\mathrm{id}_{\mathbb{Q}}\}\}}$$

and note that $U \subseteq \mathrm{Pol}(\mathfrak{D};\leq_1,\neq_1)$ since $f \in \mathrm{Pol}(\mathfrak{D})$ and $\theta_1(f)$ preserves $\leq$ and $\neq$. We claim that U contains an operation g such that $\theta_1(g)$ is dominated by the first argument. To see this, let $S \subseteq \mathbb{Q}$

be finite. If f is an ll-operation, then choose $\alpha_S \in \mathrm{Aut}(\mathbb{Q}; <)$ so that $\alpha_S(x) < 0$ for every $x \in S$. Note that $\theta_1(f((\alpha_S, \mathrm{id}_{\mathbb{Q}}), \mathrm{id}_{\mathbb{Q}^2}))$ is then dominated by the first argument on S. If $T \subseteq S$, then the homogeneity of $(\mathbb{Q}; <)$ implies that we can choose $\beta_S, \beta_T \in \mathrm{Aut}(\mathbb{Q}; <)$ such that

$$((\beta_S, \mathrm{id}_{\mathbb{Q}})f((\alpha_S, \mathrm{id}_{\mathbb{Q}}), \mathrm{id}_{\mathbb{Q}^2}))\,|_{S^2}$$

is an extension of

$$((\beta_T, \mathrm{id}_{\mathbb{Q}})f((\alpha_T, \mathrm{id}_{\mathbb{Q}}), \mathrm{id}_{\mathbb{Q}^2}))\,|_{T^2}.$$

Hence, U contains an operation g such that $\theta_1(g)$ is dominated by the first argument. Moreover, $\theta_2(g) = \theta_2(f)$ since we have only applied automorphisms that fix the second dimension, so g satisfies the statement of the lemma.

If $\theta_1(f)$ is a pp-operation, then we proceed in the same way but in this case the operation $\theta_1(g)$ is essentially unary, and by applying automorphisms and using the fact that $\mathrm{Pol}(\mathfrak{D})$ is closed we may suppose that $\theta_1(g)$ equals π_1^2.

The proof when $i = 2$ only requires flipping the inequalities in the definition of the automorphisms α_S. ∎

Corollary 4.19. *Let $f \in \mathrm{Pol}(\mathfrak{D})$ be such that $\theta_1(f)$ is an* ll- *or a* pp-*operation. Then there is an operation $g \in \mathrm{Pol}(\mathfrak{D}; \leq_1, \neq_1, \leq_2, \neq_2)$ such that $\theta_i(g)$ is dominated by the i-th argument for both $i = 1, 2$. If $\theta_1(f)$ is a* pp-*operation, then we can choose g such that $\theta_1(g) = \pi_1^2$.*

Proof. Lemma 4.17 implies that there is an operation $f' \in \mathrm{Pol}(\mathfrak{D}; \leq_2, \neq_2)$ and an index $j \in \{1, 2\}$ such that $\theta_1(f')$ is an ll-operation or a pp-operation, and $\theta_2(f')$ is dominated by the j-th argument. Hence, by Lemma 4.18 applied on f', there is $g \in \mathrm{Pol}(\mathfrak{D}; \leq_1, \neq_1, \leq_2, \neq_2)$ such that $\theta_i(g)$ is dominated by the i-th argument for both i or by the $(3-i)$-th argument for both i. Without loss of generality, we can assume that g satisfies the former, because otherwise we can replace g by the operation obtained from g by flipping arguments. Similarly, if $\theta_1(f)$ is a pp-operation, then we can choose g such that $\theta_1(g)$ is the projection π_1^2. ∎

We conclude this section with a duality result that reduces the number of cases that we have to consider in some of the forthcoming proofs.

Lemma 4.20. *Let $\mathfrak{D}$ be a first-order expansion of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. The map given by $(x, y) \mapsto (x, -y)$ is an isomorphism between $\mathfrak{D}$ and a structure which is primitively positively interdefinable with a first-order expansion $\mathfrak{C}$ of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. For every $f \in \mathrm{Pol}(\mathfrak{D})$ there exists $f' \in \mathrm{Pol}(\mathfrak{C})$ such that*

- *$\theta_1(f') = \theta_1(f)$ and*
- *$\theta_2(f') = \theta_2(f)^*$.*

Proof. For each relation R of $\mathfrak{D}$, let ϕ be the defining formula of R over $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. Replace each atomic formula of the form $x <_2 y$ in ϕ by the formula $y <_2 x$. The relation defined by the formula will be denoted by R'. The structure $\mathfrak{C}$ is then the first-order expansion of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ by all relations of the form R' where R is a relation of $\mathfrak{D}$. Then, $(x, y) \mapsto (x, -y)$ is an isomorphism between $\mathfrak{D}$ and a structure $\mathfrak{C}'$ which is the first-order expansion of $(\mathbb{Q}^2; <_1, =_1, >_2, =_2)$ by the relations R'. Note that $\mathrm{Pol}(\mathfrak{C}') = \mathrm{Pol}(\mathfrak{C})$. Let $f \in \mathrm{Pol}(\mathfrak{D})$ be of arity k. Then, the operation f' defined as

$$f'((x_1, y_1), \ldots, (x_k, y_k)) \coloneqq (\theta_1(f)(x_1, \ldots, x_k), \theta_2(f)^*(y_1, \ldots, y_k))$$

is a polymorphism of $\mathfrak{C}$ that satisfies the requirements of the lemma. ∎

4.3. Syntactic Normal Forms. In this section we prove that if $\theta_1(\mathrm{Pol}(\mathfrak{D}))$ and $\theta_2(\mathrm{Pol}(\mathfrak{D}))$ contain certain polymorphisms, then the relations of $\mathfrak{D}$ can be defined by formulas satisfying simple syntactic restrictions. These restrictions are all of the same kind: the relations can be defined by conjunctions of clauses with straightforward definitions. We start with the case that there exists an $f \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_1(f)$ is an ll-operation or a pp-operation (Proposition 4.26). We then prove a stronger statement if $\theta_1(f)$ is a pp-operation (Proposition 4.27), and an even stronger result if additionally there is no binary operation $g \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_2(g)$ is a lex-operation (Proposition 4.31). Finally, we treat the situation that for both $i \in \{1, 2\}$ there exists $f_i \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_i(f_i)$ is an ll-operation (Proposition 4.32). These results are collected in Section 4.3.2. The proofs of Propositions 4.26 and 4.27 are based on a particular normalisation of formulas that we describe in Section 4.3.1.

4.3.1. Normalisation. We will now describe a normalisation process for formulas that will be extensively used in Section 4.3.2. Let ϕ be a quantifier-free formula over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$. We may assume that

- R1. ϕ is in reduced CNF (as defined in Section 3.2),
- R2. ϕ does not contain literals of the form $x \neq y$, because such literals can be replaced by $x <_1 y \vee y <_1 x \vee x <_2 y \vee y <_2 x$, and
- R3. ϕ does not contain literals of the form $x = y$, because such literals can be replaced by $x =_1 y \wedge x =_2 y$.
- R4. ϕ does not contain literals of the form $\neg(x <_i y)$, for $i \in \{1, 2\}$, because such literals can be replaced by $y <_i x \vee y =_i x$.

We introduce two rewriting rules R5 and R6, each of which yields a formula equivalent to the original formula ϕ. The basis of both rules is the same and the only difference is in the relation contained in one of the affected literals. Suppose that ϕ contains, for distinct $i, j \in \{1, 2\}$, a clause α of the form $(u \circ_i v \vee x <_j y \vee \beta)$ where u, v, x, y are (not necessarily distinct) variables, $\circ_i \in \{<_i, =_i, \neq_i\}$ and let ϕ' be the other clauses of ϕ. If

$$\phi' \wedge \neg\beta \wedge u \circ_i v \text{ implies } x =_j y,$$

then we replace α by the two clauses

$$\begin{aligned} &(x <_j y \vee x =_j y \vee \beta) \\ \text{and } &(u \circ_i v \vee x \neq_j y \vee \beta). \end{aligned}$$

If the relation $\circ_i$ is $<_i$, then we will refer to the rewriting rule as R5, and otherwise (that is, when $\circ_i \in \{=_i, \neq_i\}$) we will refer to the rule as R6. To see that the new formula is equivalent to ϕ, let s be a solution to ϕ. If s satisfies β, then the two new clauses are satisfied. If s does not satisfy β, then it must satisfy $u \circ_i v$ or $x <_j y$. In the first case, s satisfies the second new clause, and by assumption it also satisfies the first new clause. In the latter case, it clearly satisfies both the first and the second clause. Now suppose that conversely, s satisfies ϕ' and the two new clauses. If s satisfies β or $x <_j y$ then α is satisfied. Otherwise, the first new clause implies that $x =_j y$, and hence the second clause implies that $u \circ_i v$, and hence α is satisfied. If the formula obtained from applying R5 or R6 is not reduced, we remove literals to make it reduced. Note that after every application of R5 or R6 the conditions R1–R4 will still be satisfied.

Example 4.21. We give an example of the rewriting procedure starting from a formula ϕ in reduced CNF.

$$(x =_1 y) \wedge (u =_1 v \vee x = y) \wedge (u <_1 v \vee x <_2 y)$$
$$(x =_1 y) \wedge (u =_1 v \vee (x =_1 y \wedge x =_2 y)) \wedge (u <_1 v \vee x <_2 y) \quad \text{(R3)}$$
$$(x =_1 y) \wedge (u =_1 v \vee x =_1 y) \wedge (u =_1 v \vee x =_2 y) \wedge (u <_1 v \vee x <_2 y) \quad \text{(R1)}$$
$$(x =_1 y) \wedge (u =_1 v \vee x =_2 y) \wedge (u <_1 v \vee x <_2 y) \quad \text{(R1)}$$
$$(x =_1 y) \wedge (u =_1 v \vee x =_2 y) \wedge (x <_2 y \vee x =_2 y) \wedge (u <_1 v \vee x \neq_2 y) \quad \text{(R5)}$$
$$(x =_1 y) \wedge (u =_1 v \vee x =_2 y) \wedge (x <_2 y \vee x =_2 y) \wedge (u <_1 v \vee u =_1 v) \wedge (x \neq_2 y \vee u \neq_1 v) \quad \text{(R6)}$$

The resulting formula does not admit application of any of the rewriting rules.

The reason to split the rewriting rule into two rules R5 and R6 is that only R5 terminates (i.e. can be applied only finitely many times) on every quantifier-free CNF formula (Lemma 4.22). To prove termination of R6 (Lemma 4.24) and existence of an equivalent formula to which none of the rewriting rules above can be applied, we require the existence of a certain operation that preserves the formula. On the way we also prove a syntactic restriction on such formulas (Lemma 4.23). Note that Lemma 4.23 and 4.24 remain true also if $\theta_1(f)$ is the dual of an ll- or pp-operation; it can be proved using the version of Corollary 4.19 based on duals.

LEMMA 4.22. *Let ϕ be a quantifier-free CNF formula over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$. Then the rewriting rule R5 applied on ϕ terminates.*

Proof. Let α be an arbitrary clause over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$. Let $m(\alpha)$ denote the total number of $<_1$- and $<_2$-literals in α. Assume we apply R5 to the clause $\alpha = (u <_i v \vee x <_j y \vee \beta)$ where we assume (without loss of generality) that $u <_i v \notin \beta$ and $x <_j y \notin \beta$. This yields two clauses $\alpha_1 = (x <_j y \vee x =_j y \vee \beta)$ and $\alpha_2 = (u <_i v \vee x \neq_j y \vee \beta)$. Note that $m(\alpha_1) < m(\alpha)$ and $m(\alpha_2) < m(\alpha)$ and reducing the formula cannot increase $m(\alpha)$ for any clause α.

Now consider a quantifier-free CNF formula ϕ over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$. Let ℓ be the maximum clause length of ϕ and let k be the number of variables appearing in ϕ. Note that R5-rewriting cannot increase ℓ or k, and for any clause α in ϕ, it holds that $0 \leq m(\alpha) \leq \ell$. Let $f(\ell, k)$ denote the (finite) number of possible clauses over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ where clause length is bounded by ℓ and at most k variables are used. Arbitrarily choose a clause α in ϕ. If R5 is applied to α, then we know that α is replaced by at most two new clauses α_1 and α_2 where $m(\alpha_1) < m(\alpha)$ and $m(\alpha_2) < m(\alpha)$. Thus, the clause α can result in at most 2^ℓ applications of rule R5. Since there are at most $f(\ell, k)$ clauses in ϕ, we conclude that R5 can be applied at most $f(\ell, k) \cdot 2^\ell$ times to ϕ. ☐

We continue with rewriting rule R6. The following notation will be practical in several proofs dealing with syntactic forms. Let g be a binary operation on $\mathbb{Q}$, ϕ a formula over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ and set of variables X and $s, t : X \to \mathbb{Q}$ assignments of ϕ. Then $g(s, t)$ represents the assignment of ϕ that assigns to variable $x \in X$ the value $g(s(x), t(x))$.

LEMMA 4.23. *Let ϕ be a quantifier-free CNF formula over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ such that R5 cannot be applied to ϕ. Suppose that ϕ is preserved by an operation $f \in \mathrm{Pol}(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ such that $\theta_1(f)$ is an* ll*-operation or a* pp*-operation. Then ϕ does not contain a clause that contains a $<_j$-literal for both $j \in \{1, 2\}$.*

Proof. By Corollary 4.19, there is $g \in \mathrm{Pol}(\mathbb{Q}^2; =_1, <_1, =_2, <_2, \leq_1, \neq_1, \leq_2, \neq_2)$ that preserves ϕ such that $\theta_j(g)$ is dominated by j-th argument for both j.

Suppose for contradiction that ϕ contains a clause α of the form $(u <_1 v \vee x <_2 y \vee \beta)$. Since R5 cannot be applied and ϕ is reduced, there are satisfying assignments s and t of ϕ such that s

satisfies $u <_1 v$, $y <_2 x$, and falsifies β and t satisfies $v <_1 u$, $x <_2 y$, and falsifies β. Then the tuple $g(t,s)$ satisfies $v <_1 u$ since $\theta_1(g)$ is dominated by the first argument, and it satisfies $y <_2 x$ since $\theta_2(g)$ is dominated by the second argument. Moreover, all other literals of α are falsified, too, since g preserves $<_j$, $=_j$, $\leq_j$, and $\neq_j$ for $j \in \{1,2\}$. ☐

LEMMA 4.24. *Let ϕ be a quantifier-free CNF formula over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ such that R5 cannot be applied on ϕ. Suppose that ϕ is preserved by an operation $f \in \mathrm{Pol}(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ such that $\theta_1(f)$ is an* ll*-operation or a* pp*-operation. Then the rewriting rule R6 applied to ϕ terminates.*

Proof. By Lemma 4.23, ϕ does not contain a clause that contains a $<_j$-literal for both $j \in \{1,2\}$. Note that by an application of the rewriting rule R6 no such clause can be created. For a clause α in ϕ, let $p(\alpha)$ be the number of pairs of a $\{=_i, \neq_i\}$-literal and a $<_j$-literal, for distinct i and j, that appears in α.

Arbitrarily choose a clause α in ϕ that admits an application of R6. Then $\alpha = (u \circ_i v \vee x <_j y \vee \beta)$ for $\circ_i \in \{=_i, \neq_i\}$. After applying R6, α is replaced by two clauses $\alpha_1 = (x <_j y \vee x =_j y \vee \beta)$ and $\alpha_2 = (u \circ_i v \vee x \neq_j y \vee \beta)$. Observe that $p(\alpha_1) < p(\alpha)$ and $p(\alpha_2) < p(\alpha)$ since α does not contain $<_{3-j}$-literals. Moreover, reducing the formula does not increase $p(\gamma)$ for any clause γ in the formula.

Let ℓ be the maximum clause length of ϕ. For any clause γ in ϕ, it holds that $0 \leq p(\gamma) \leq \ell^2/4$. Using an argument analogous to the one in Lemma 4.22, we conclude that R6 can be applied only finitely many times to ϕ. ☐

If ϕ is a reduced formula such that none of the rewriting rules presented above are applicable, then we call it *normal*. If ϕ is a quantifier-free CNF formula over $(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ preserved by $f \in \mathrm{Pol}(\mathbb{Q}^2; =_1, <_1, =_2, <_2)$ where $\theta_1(f)$ is an ll-operation, a pp-operation, or the dual of such an operation, it is possible to rewrite it to an equivalent normal formula by first applying R5 until it terminates and then applying R6 until it terminates. Observe that the resulting formula satisfies conditions R1–R4 and does not admit an application of R5 (since it does not contain a clause of the required form by Lemma 4.23) or R6. Note that if ϕ is normal and contains a clause of the form $(u \circ_i v \ \vee \ x <_j y \ \vee \ \beta)$, then ϕ has a satisfying assignment that satisfies $u \circ_i v$, falsifies β, and satisfies $y <_j x$, because otherwise we could have applied R5 or R6.

The rewriting rules and Lemma 4.22 can be generalised in a straightforward fashion to formulas over $(\mathbb{Q}; <)^{(n)}$. To prove a generalised version of Lemma 4.23, one needs to use Corollary 4.19 to produce a polymorphism with the particular domination property in two distinct dimensions i and j (see the proof of Proposition 4.37 for more details). The generalised version of the lemma then shows that there is no clause containing both $<_i$-literals and $<_j$-literals for distinct i and j under the assumption that for all but at most one $p \in \{1, \ldots, n\}$ there is $f_p \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(f_p)$ is an ll-operation or a pp-operation. Subsequently, a generalised version of Lemma 4.24 can be proved and normal formulas can be defined. The normalisation process for formulas over $(\mathbb{Q}; <)^{(n)}$ will be used in Section 4.5.

4.3.2. Definitions via Restricted Clauses. The following definition is central in our presentation of various syntactic normal forms.

DEFINITION 4.25. *A clause is* weakly 1-determined *if it is of the form*

$$\psi \vee \bigvee_{i \in \{1,\ldots,k\}} x_i \neq_2 y_i$$

where ψ is 1-determined and $k \geq 0$. Weakly 2-determined clauses are defined analogously.

A clause can simultaneously be weakly 1-determined *and* weakly 2-determined: $x \neq_1 y \vee u \neq_2 v$ is one example. Normalised formulas in the sense of Section 4.3.1 play a key role in our first result concerning logical definitions based on weakly i-determined clauses.

PROPOSITION 4.26. *Suppose that* $\mathrm{Pol}(\mathfrak{D})$ *contains an operation* f *such that* $\theta_1(f)$ *is an* ll-*operation or a* pp-*operation. Then, the following holds for every relation* R *in* $\mathfrak{D}$*: if a normal formula* ϕ *is a definition of* R *over* $(\mathbb{Q}^2, <_1, =_1, <_2, =_2)$*, then* ϕ *is a conjunction of clauses each of which is weakly* i*-determined for some* $i \in \{1, 2\}$*.*

Proof. By Corollary 4.19, there is an operation $g \in \mathrm{Pol}(\mathfrak{D}; \leq_1, \neq_1, \leq_2, \neq_2)$ such that $\theta_i(g)$ is dominated by the i-th argument for both $i \in \{1, 2\}$. Let ϕ be a normal formula that defines a relation $R \in \mathfrak{D}$ over $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. Note, in particular, that ϕ cannot be rewritten using rule R5 or R6. Let ψ be a clause of ϕ. Properties R1–R4 imply that ψ is equal to

$$\psi_1 \vee \psi_2 \vee \bigvee_{l \in \{1,\ldots,k\} \text{ and } j \in \{1,2\}} x_l \neq_j y_l$$

where ψ_i for $i \in \{1, 2\}$ only contains literals of the form $x =_i y$ or $x <_i y$.

We show the result by verifying that ψ_1 or ψ_2 is the empty disjunction. Suppose to the contrary that ψ_1 contains a literal ℓ_1 and ψ_2 contains a literal ℓ_2. Since ϕ is reduced, it must have a satisfying assignment s that satisfies ℓ_1 and falsifies all other literals of ψ, and there also exists a satisfying assignment t that satisfies ℓ_2 and falsifies all other literals of ψ. Note that by Lemma 4.23 it cannot occur that ℓ_1 equals $u <_1 v$ and ℓ_2 equals $x <_2 y$, since R5 cannot be applied to ϕ. Therefore, we have to consider the following cases.

1. Suppose that the literal ℓ_1 is of the form $u =_1 v$ and the literal ℓ_2 equals $x <_2 y$. Since ϕ is normal, we may assume (by R6) that s satisfies $y <_2 x$. Then $g(t, s)$ satisfies $u \neq_1 v$ since $\theta_1(g)$ is dominated by the first argument, and it satisfies $y <_2 x$ because $\theta_2(g)$ is dominated by the second argument. Moreover, all other literals of ψ are falsified. Hence, $g(t, s)$ does not satisfy ϕ, which contradicts $g \in \mathrm{Pol}(\mathfrak{D})$.
2. The case that the literal ℓ_1 equals $u <_1 v$ and the literal ℓ_2 equals $x =_2 y$, and the case that ℓ_1 equals $u =_1 v$ and the literal ℓ_2 equals $x =_2 y$ can be treated similarly.

If ψ_1 is empty, then we obtain a clause that is weakly 2-determined. Likewise, if ψ_2 is empty, then ψ is a weakly 1-determined clause. ∎

Under additional conditions on polymorphisms, we can define relations by formulas that are based on weakly 1-determined clauses together with (not weakly) 2-determined clauses.

PROPOSITION 4.27. *Let* $f \in \mathrm{Pol}(\mathfrak{D})$ *be such that* $\theta_1(f)$ *is a* pp-*operation. Then every normal conjunction* ϕ *of weakly 1-determined and weakly 2-determined clauses that is preserved by* f *is a conjunction of weakly* 1*-determined and of* 2*-determined clauses.*

Proof. By Corollary 4.19 there is an operation $g \in \mathrm{Pol}(\mathfrak{D}; \leq_1, \neq_1, \leq_2, \neq_2)$ such that $\theta_1(g) = \pi^2_1$ and $\theta_2(g)$ is dominated by the second argument. Suppose for contradiction that ϕ contains a clause ψ with a literal $x \neq_1 y$ and a literal χ which is of the form $u <_2 v$ or $u =_2 v$. The formula ϕ is reduced so it has a satisfying assignment s which satisfies $x \neq_1 y$ and falsifies all other literals of ψ, and a satisfying assignment t which satisfies χ and falsifies all other literals in ψ; in particular, t satisfies $x =_1 y$.

We first consider the case that χ is of the form $u <_2 v$ and s consequently satisfies $v \leq_2 u$. Since ϕ is normal we may even suppose that s satisfies $v <_2 u$. Then $g(t, s)$ satisfies $x =_1 y$ since $\theta_1(g) = \pi^2_1$ and it satisfies $v <_2 u$ since $\theta_2(g)$ is dominated by the second argument. Hence, it satisfies neither the literal $x \neq_1 y$ nor the literal $u <_2 v$, nor any of the other literals of ψ since f preserves $\leq_i$ and $\neq_i$ for $i \in \{1, 2\}$. This is in contradiction to the assumption that ϕ is preserved by g.

The case that χ is of the form $u =_2 v$ similarly leads to a contradiction. □

Remark 4.28. Note that it is not true that every weakly 2-determined clause of formula ϕ in Proposition 4.27 is 2-determined: For example, consider the clause ψ of the form $x \neq_1 y \vee u \neq_2 v$. The clause ψ is weakly 2-determined, but not 2-determined, and it satisfies the assumptions of Proposition 4.27: it is normal, and preserved by a map (f_1, f_2) where f_1 is a pp-operation and f_2 is an ll-operation. To see this, let s, t be two satisfying assignments of ψ. Either one of s and t satisfies $u \neq_2 v$ and hence $(f_1, f_2)(s, t)$ satisfies it as well by the injectivity of f_2, or both s and t satisfy $x \neq_1 y$ and hence $(f_1, f_2)(s, t)$ satisfies it as well since f_1 preserves $\neq$.

In the next proof we use the notion of the orbit of a k-tuple $(t_1, \dots, t_k) \in \mathbb{Q}^k$ under $\mathrm{Aut}(\mathbb{Q}; <)$ from Section 2.3. Observe that the homogeneity of $(\mathbb{Q}; <)$ implies that the orbit of a tuple $(t_1, \dots, t_k)$ under $\mathrm{Aut}(\mathbb{Q}; <)$ is determined by the weak linear order induced on $(t_1, \dots, t_k)$ in $(\mathbb{Q}; <)$. We need a weak linear order since some of the elements $t_1, \dots, t_k$ may be equal.

PROPOSITION 4.29. *As in the previous proposition, let $f \in \mathrm{Pol}(\mathfrak{D})$ be such that $\theta_1(f)$ is a* pp-*operation. Then every relation of $\mathfrak{D}$ can be defined by a conjunction of* 2*-determined clauses and weakly* 1*-determined clauses of the form*

$$u_1 \neq_2 v_1 \vee \cdots \vee u_m \neq_2 v_m \vee y_1 \neq_1 x \vee \cdots \vee y_k \neq_1 x \vee z_1 \leq_1 x \vee \cdots \vee z_l \leq_1 x. \tag{4.1}$$

(In other words, if we drop the first m literals in and remove subscripts we obtain a formula as described in Theorem 4.7.)

Proof. Let ϕ be a formula over a finite set of variables X that defines a relation R from $\mathfrak{D}$. Without loss of generality, we may assume that ϕ is normal. Hence, by Proposition 4.26 and Proposition 4.27, ϕ is a conjunction of weakly 1-determined and of 2-determined clauses. Let ϕ' be the conjunction of all clauses of the form (4.1) and of all 2-determined clauses with variables from X that are reduced and implied by ϕ; since $|X|$ is finite, ϕ' is a finite formula. We claim that ϕ' implies ϕ, and consequently that ϕ' is a definition of R of the required syntactic form.

The 2-determined conjuncts of ϕ are clearly implied by ϕ'. In the rest of the proof, we prove in two steps that every weakly 1-determined conjunct of ϕ is implied by ϕ': Firstly, we show that we may assume that all weakly 1-determined clauses of ϕ are of the form (4.2) and that they are minimal in a particular sense specified below. Secondly, we show that every such clause is implied by ϕ', because the assumption that R is preserved by f would be violated otherwise.

To proceed with the first step, we claim that every conjunction of weakly 1-determined clauses can be written as a conjunction of formulas of the form

$$\chi \vee \neg(y_1 \circ_1 y_2 \wedge \cdots \wedge y_k \circ_k y_{k+1})$$

where $\chi := \bigvee_{i=1}^m u_i \neq_2 v_i$ with $u_1, \dots, u_m,\ v_1, \dots, v_m,\ y_1, \dots, y_{k+1}$ being variables from X, and $\circ_1, \dots, \circ_k \in \{=_1, <_1\}$. To see this, first note that every orbit of $k+1$-tuples in $\mathrm{Aut}(\mathbb{Q}; <)$ can be defined by a formula of the form $x_1 \circ_1 x_2 \wedge \cdots \wedge x_k \circ_k x_{k+1}$ where $\circ_1, \dots, \circ_k \in \{=, <\}$ if the variables are named appropriately. Hence, every first-order formula in $(\mathbb{Q}; <)$ is equivalent to a conjunction of negations of such formulas. It follows that every 1-determined clause can be written as a conjunction of formulas of the form $\neg(y_1 \circ_1 y_2 \wedge \cdots \wedge y_k \circ_k y_{k+1})$ for $\circ_1, \dots, \circ_k \in \{=_1, <_1\}$. Using distributivity of disjunction over conjunction we can therefore rewrite a conjunction of weakly 1-determined clauses into a conjunction of formulas of the desired form.

We may henceforth assume that every conjunct ψ of ϕ that is not 2-determined is of the form

$$\chi \vee y_1 \circ_1 y_2 \vee \cdots \vee y_k \circ_k y_{k+1} \tag{4.2}$$

where χ is as above and $\circ_1, \ldots, \circ_k \in \{\neq_1, \geq_1\}$. We may additionally assume that ϕ contains only those clauses of the form (4.2) that are minimal in the following sense: a clause $\psi = \chi \vee y_1 \circ_1 y_2 \vee \cdots \vee y_k \circ_k y_{k+1}$ is minimal if there is no clause $\chi \vee y'_1 \circ_1 y'_2 \vee \cdots \vee y'_{k'} \circ_{k'} y'_{k'+1}$ implied by ϕ such that it implies ψ, it is different from ψ and $y'_1, \ldots, y'_{k'}$ is a subsequence of $y_1, \ldots, y_k$.

We now show that the clause ψ is implied by ϕ'. If $\circ_i$ equals $\neq_1$ for every $i \in \{1, \ldots, k\}$ then ψ is equivalent to $\chi \vee \neg(y_1 =_1 y_{k+1} \wedge \cdots \wedge y_k =_1 y_{k+1})$ and hence is equivalent to $\chi \vee y_1 \neq_1 y_{k+1} \vee \cdots \vee y_k \neq_1 y_{k+1}$ which is of the form (4.1) (for $l = 0$). But this formula is then a conjunct of ϕ' and there is nothing to be shown.

Otherwise, let j be smallest such that $\circ_j$ equals $\geq_1$. Then ψ is equivalent to a formula of the form

$$\chi \vee y_1 \neq_1 y_j \vee \cdots \vee y_{j-1} \neq_1 y_j \vee y_j \geq_1 y_{j+1} \vee \eta$$

where η is of the form $y_{j+1} \circ_{j+1} y_{j+2} \vee \cdots \vee y_k \circ_k y_{k+1}$ for $\circ_{j+1}, \ldots, \circ_k \in \{\neq_1, \geq_1\}$. The formula

$$\chi \vee y_1 \neq_1 y_j \vee \cdots \vee y_{j-1} \neq_1 y_j \vee y_{j+1} \leq_1 y_j \vee \cdots \vee y_{k+1} \leq_1 y_j \tag{4.3}$$

implies ψ. To see this, compare the negations of the formulas: $\neg\psi$ is equivalent to the conjunction of $\neg\chi \wedge y_1 =_1 \cdots =_1 y_j \wedge y_j <_1 y_{j+1}$ and the fact that $y_{j+1}, \ldots, y_k$ is a non-decreasing sequence, while the negation of the formula (4.3) is equivalent to

$$\neg\chi \wedge y_1 =_1 \cdots =_1 y_j \wedge y_{j+1} >_1 y_j \wedge \cdots \wedge y_k >_1 y_j.$$

Therefore, if the formula (4.3) is a conjunct of ϕ', then there is again nothing to be shown.

Otherwise, there must exist an assignment r that satisfies ϕ but not (4.3). Note that $r(y_j) <_1 r(y_i)$ for every $i \in \{j+1, \ldots, k+1\}$. Since ψ was minimal, the clause $\chi \vee \eta$ is not implied by ϕ and hence there must also exist an assignment s which satisfies ϕ, but does not satisfy $\chi \vee \eta$. Choose $\alpha \in \mathrm{Aut}(\mathfrak{D})$ such that $\alpha(r(y_j)) <_1 0 <_1 \alpha(r(y_i))$ for all $i \in \{j+1, \ldots, k+1\}$. Then $t := f(\alpha(r), s)$ is an assignment that does not satisfy ψ by the definition of a pp-operation. This contradicts that ϕ defines a relation from $\mathfrak{D}$. ∎

As the following lemma shows, when $\theta_1(\mathrm{Pol}(\mathfrak{D}))$ contains a pp-operation and $\theta_2(\mathrm{Pol}(\mathfrak{D}))$ contains an ll-operation, we may restrict to formulas of a very particular form.

Proposition 4.30. *Let $f_1, f_2 \in \mathrm{Pol}(\mathfrak{D})$ be such that $\theta_1(f_1)$ is a* pp*-operation and $\theta_2(f_2)$ is an* ll*-operation. Then every relation of $\mathfrak{D}$ can be defined by a conjunction of weakly* 1*-determined clauses of the form (4.1) and* 2*-determined clauses*

$$x_1 \neq_2 y_1 \vee \cdots \vee x_m \neq_2 y_m \vee z_1 <_2 z_0 \vee \cdots \vee z_\ell <_2 z_0 \vee (z_0 =_2 z_1 =_2 \cdots =_2 z_\ell), \tag{4.4}$$

where it is permitted that $l = 0$ or $m = 0$, and the final disjunct may be omitted (in other words, clauses obtained from ll-Horn clauses by adding the subscript 2 *to all relation symbols).*

Proof. Let R be a relation of $\mathfrak{D}$. By Proposition 4.29, R can be defined by a formula $\phi_1 \wedge \phi_2$, where ϕ_1 is a conjunction of weakly 1-determined clauses of the form (4.1) and ϕ_2 is a conjunction of 2-determined clauses. Recall the operator cr introduced at the end of Section 3.2. Let us denote $\mathrm{cr}(\phi_1 \wedge \phi_2, \{2\}, \phi_2)$ by ψ_2; note that ψ_2 is a conjunction of 2-determined clauses preserved by $\mathrm{Pol}(\mathfrak{D})$. Moreover the formula $\phi_1 \wedge \psi_2$ still defines R. By Lemma 3.11, $\hat{\psi}_2$ is preserved by an ll-operation. By Theorem 4.12, $\hat{\psi}_2$ may be taken to be a conjunction of ll-Horn clauses and therefore ψ_2 would be a conjunction of clauses of the form (4.4). This concludes the proof. ∎

We next present an even more restricted syntactic form; in this case it is sufficient to use i-determined clauses, $i \in \{1,2\}$, and we do not need weakly i-determined clauses at all.

PROPOSITION 4.31. *Suppose that there exists $f \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_1(f)$ is a* pp*-operation, but there is no binary $g \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_2(g)$ is a* lex*-operation. Then every relation of $\mathfrak{D}$ can be defined by a formula $\phi_1 \wedge \phi_2$ such that ϕ_i is a conjunction of i-determined clauses for $i = 1, 2$. Moreover, for every such definition $\phi_1 \wedge \phi_2$, there is a conjunction ψ_1 of 1-determined clauses of the form*

$$y_1 \neq_1 x \vee \cdots \vee y_k \neq_1 x \vee z_1 \leq_1 x \vee \cdots \vee z_l \leq_1 x$$

such that $\psi_1 \wedge \phi_2$ still defines the same relation.

Proof. Lemma 4.17 implies that $\mathrm{Pol}(\mathfrak{D})$ contains an operation f' such that $\theta_1(f')$ is a pp-operation and $\theta_2(f')$ is either a lex-operation or it is essentially unary; by assumption, it cannot be a lex-operation so it must be essentially unary. Let $i \in \{1,2\}$ be such that $\theta_2(f')$ depends only on the i-th argument. Since $\theta_2(f')$ preserves $<$, the set $\overline{\{(\mathrm{id}_{\mathbb{Q}}, \alpha) f' \mid \alpha \in \mathrm{Aut}(\mathbb{Q}; <)\}} \subseteq \mathrm{Pol}(\mathfrak{D})$ contains an operation f'' such that $\theta_1(f'') = \theta_1(f')$ and $\theta_2(f'') = \pi^2_i$. Therefore, we can assume without loss of generality that $\theta_2(f') = \pi^2_i$. By Lemma 4.18 applied on f', we obtain $g \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_1(g) = \pi^2_{3-i}$ and $\theta_2(g) = \pi^2_i$. This implies that $(\pi^2_1, \pi^2_2) \in \mathrm{Pol}(\mathfrak{D})$. By Lemma 3.12, every relation of $\mathfrak{D}$ can be defined by a conjunction of 1-determined and 2-determined clauses.

Let R be a relation of $\mathfrak{D}$. Let $\phi_1 \wedge \phi_2$ be a definition of R such that ϕ_i is a conjunction of i-determined clauses, $i = 1, 2$. Recall the operator $\mathrm{cr}(\cdot)$ introduced at the end of Section 3.2. Let $\psi_1 = \mathrm{cr}(\phi_1 \wedge \phi_2, \{1\}, \phi_1)$. Then the formula ψ_1 is a conjunction of 1-determined clauses preserved by $\mathrm{Pol}(\mathfrak{D})$ and $\psi_1 \wedge \phi_2$ defines R. By Lemma 3.11, $\hat{\psi_1}$ is preserved by a pp-operation. Hence, by Theorem 4.7, we may assume that ψ_1 is a conjunction of clauses of the form

$$y_1 \neq_1 x \vee \cdots \vee y_k \neq_1 x \vee z_1 \leq_1 x \vee \cdots \vee z_l \leq_1 x.$$

This concludes the proof. ∎

Assume that for every $i \in \{1,2\}$, the clone $\mathrm{Pol}(\mathfrak{D})$ contains an operation f_i such that $\theta_i(f_i)$ is an ll-operation or the dual of such an operation. Then, we may combine the information about syntactically restricted definitions of the relations of $\mathfrak{D}$ from Proposition 4.26 with ll-Horn definability from Theorem 4.12, and obtain the next result (Proposition 4.32). We will use the following notation for simplifying the presentation. For a formula ϕ over $(\mathbb{Q}^n; <_1, =_1, \ldots, <_n, =_n)$, we introduce n fresh variables $x^1, \ldots, x^n$ for each variable x that appears in ϕ. Then, we let $\mathrm{ve}(\phi)$ (for *variable expansion*) denote the formula over $(\mathbb{Q}; <)$ resulting from ϕ by replacing each atomic formula of the form $x \circ_i y$ by $x^i \circ y^i$, where $\circ \in \{<, \leq, =, \neq\}$ and $i \in \{1, \ldots, n\}$.

PROPOSITION 4.32. *Suppose that for every $i \in \{1,2\}$ the clone $\mathrm{Pol}(\mathfrak{D})$ contains an operation f_i such that $\theta_i(f_i)$ is an* ll*-operation. Then every relation of $\mathfrak{D}$ has a definition by a conjunction of clauses of the form*

$$x_1 \neq_{i_1} y_1 \vee \cdots \vee x_m \neq_{i_m} y_m \vee z_1 <_j z_0 \vee \cdots \vee z_\ell <_j z_0 \vee (z_0 =_j z_1 =_j \cdots =_j z_\ell)$$

for $i_1, \ldots, i_m, j \in \{1,2\}$ and where the last disjunct may be omitted. Moreover, $\mathfrak{D}$ has a primitive positive interpretation in $\mathfrak{L}$.

Proof. Let R be a relation of $\mathfrak{D}$, and let ϕ be a definition of R. Without loss of generality, we may assume that ϕ is normal and hence a conjunction of weakly 1-determined and weakly 2-determined clauses (Proposition 4.26).

Claim. The formula $\text{ve}(\phi)$ over $(\mathbb{Q}; <)$ is preserved by every ll operation.

To prove the claim, let ll be an ll-operation and let r' and s' be satisfying assignments for $\text{ve}(\phi)$. We have to show that $t'(x) := \text{ll}(r'(x), s'(x))$ satisfies $\text{ve}(\phi)$. Let ψ' be a conjunct of $\text{ve}(\phi)$. Then ψ' has been created from a conjunct ψ of ϕ with variables $y_1, \dots, y_m$, which must be weakly i-determined, for some $i \in \{1, 2\}$. We assume henceforth that $i = 1$; the other case can be shown analogously. Note that the maps $r \colon x \mapsto (r'(x^1), r'(x^2))$ and $s \colon x \mapsto (s'(x^1), s'(x^2))$ are satisfying assignments to ϕ. Since $\theta_1(f_1)$ is an ll-operation, we may choose $\alpha \in \text{Aut}(\mathbb{Q}; <)$ such that for $t(x) := (\alpha, \text{id}) f_1\big(r(x), s(x)\big)$ we have $(t(y_1)_1, \dots, t(y_m)_1) = (t'(y_1^1), \dots, t'(y_m^1))$ (see Remark 4.5). Therefore, we are done if one of the disjuncts of ψ of the form $x =_1 y$, $x <_1 y$, or $x \neq_1 y$ is satisfied by t, because then a disjunct of ψ' of the form $x^1 = y^1$, $x^1 < y^1$, or $x^1 \neq y^1$ is satisfied by t'. Otherwise, since t satisfies ψ, there must be a literal of ψ' of the form $x^2 \neq y^2$. We claim that t' satisfies this literal. As t satisfies $x \neq_2 y$, we must have $r(x) \neq_2 r(y)$ or $s(x) \neq_2 s(y)$, so $r'(x^2) \neq r'(y^2)$ or $s'(x^2) \neq s'(y^2)$. Hence, $t'(x^2) \neq t'(y^2)$ by the injectivity of ll. ⋄

Note that the first statement of the proposition follows from the claim by Theorem 4.12. The claim and Theorem 4.12 also imply that we obtain a two-dimensional primitive positive interpretation of $\mathfrak{D}$ in $\mathfrak{L}$, which proves the second statement of the theorem. □

4.4. Classification in the 2-Dimensional Case. The aim of this section is to present our algebraic dichotomy result for polymorphism clones of first-order expansions of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ together with a corresponding complexity dichotomy (Section 4.4.2) We begin by presenting algorithmic results.

4.4.1. Polynomial-time Algorithms. In this section we present polynomial-time solvability results for $\text{CSP}(\mathfrak{D})$ when $\mathfrak{D}$ is a first-order expansion of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ such that $\theta_1(\text{Pol}(\mathfrak{D}))$ and $\theta_2(\text{Pol}(\mathfrak{D}))$ contain sufficiently strong polymorphisms. Intuitively speaking, in every tractable case, there must be in each of the dimensions a polymorphism that guarantees tractability. By Theorem 4.3, this means that both $\theta_1(\text{Pol}(\mathfrak{D}))$ and $\theta_2(\text{Pol}(\mathfrak{D}))$ have to contain one of the operations min_3, mx_3, mi_3 or ll_3. The case when ll_3, equivalently any ll-operation, lies in $\theta_i(\text{Pol}(\mathfrak{D}))$ for both i is treated in Proposition 4.33. The case when $\theta_1(\text{Pol}(\mathfrak{D}))$ contains min_3, mx_3, mi_3 and $\theta_2(\text{Pol}(\mathfrak{D}))$ contains ll_3 is covered in Proposition 4.34; the case with the dimensions switched is symmetric. The case when both $\theta_1(\text{Pol}(\mathfrak{D}))$ and $\theta_2(\text{Pol}(\mathfrak{D}))$ contain one of the operations min_3, mx_3, or mi_3 can be easily handled using Proposition 4.31 and Lemma 3.12 and is treated directly in the proof of the classification theorem (Theorem 4.35).

Proposition 4.33. *Suppose that $\mathfrak{D}$ has a finite relational signature and for every $i \in \{1, 2\}$, the clone* $\text{Pol}(\mathfrak{D})$ *contains an operation f_i such that $\theta_i(f_i)$ is an* ll*-operation. Then* $\text{CSP}(\mathfrak{D})$ *can be solved in polynomial time.*

Proof. The structure $\mathfrak{D}$ has a primitive positive interpretation in $\mathfrak{L}$ by Proposition 4.32. The result follows from Lemma 2.9 since $\text{CSP}(\mathfrak{L})$ can be solved in polynomial time. □

We will use some additional observations and notations in the proof of the next proposition (and also in its generalisation Proposition 4.39). Assume that the relational structure $\mathfrak{A}$ with domain $\mathbb{Q}$ has a finite signature and there is an ll-operation in $\text{Pol}(\mathfrak{A})$. We know from Theorem 4.3 that $\text{CSP}(\mathfrak{A})$ is polynomial-time solvable. Assume now that $\mathfrak{A}'$ is a solvable instance of $\text{CSP}(\mathfrak{A})$ with variable set X. The polynomial-time algorithm for $\text{CSP}(\mathfrak{A})$ by Bodirsky and Kára [32, Section 4] computes additional information in the form of an *equality set*: a set $E \subseteq X^2$ such that for every $(x, x') \in E$ and every solution s, it holds that $s(x) = s(x')$. Furthermore, there exists a solution t such that $t(x) \neq t(x')$ for every $(x, x') \notin E$.

With this in mind, we introduce the following notation. Let X be a set of variables, $i \in \{1, 2\}$, and $E \subseteq \{(x^i, y^i) \mid x, y \in X\}$ where x^i, y^i denote fresh variables. Assume that ϕ is a formula over $(\mathbb{Q}^2, <_1, =_1, \ldots, <_n, =_n)$ that only uses variables from X. If ϕ equals $\phi_1 \wedge \phi_2$, where ϕ_i is a conjunction of i-determined clauses and ϕ_{3-i} is a conjunction of weakly $(3-i)$-determined clauses, then we let $\mathrm{cm}(\phi, E)$ (for *clause modification*) denote the formula resulting from ϕ by performing the following procedure:

- If $(x^i, y^i) \in E$ and ϕ_{3-i} contains a clause with the literal $x \neq_i y$, then remove this literal from the clause and add the conjunct $x =_i y$ to ϕ.
- If $(x^i, y^i) \notin E$ and ϕ_{3-i} contains a clause with the literal $x \neq_i y$, then delete all literals in the clause but this one.

For example, if $i = 1$, $\phi = (x =_1 y \vee x <_1 z) \wedge (x \neq_1 z \vee u <_2 v) \wedge (u \neq_1 z \vee x =_2 y)$ and $E = \{(x^1, z^1)\}$, then $\mathrm{cm}(\phi, E) = (x =_1 y \vee x <_1 z) \wedge (x =_1 z) \wedge (u <_2 v) \wedge (u \neq_1 z)$.

The following proposition and its proof describe a polynomial time algorithm for $\mathrm{CSP}(\mathfrak{D})$, where $\mathrm{Pol}(\mathfrak{D})$ satisfies certain conditions (see Algorithm 4.1). A generalised version of this algorithm for first-order expansions of $(\mathbb{Q}, <)^{(n)}$ with $n \geq 2$ will be presented in Proposition 4.39 and Algorithm 4.2. There are no profound differences between the algorithms but the case when $n = 2$ is easier to present since we can keep the formal machinery at a minimum.

Proposition 4.34. *Suppose that $\mathfrak{D}$ has a finite relational signature and that $\mathrm{Pol}(\mathfrak{D})$ contains f_1, f_2 such that $\theta_1(f_1)$ equals min_3, mx_3, or mi_3, and $\theta_2(f_2)$ is an ll-operation. Then $\mathrm{CSP}(\mathfrak{D})$ can be solved in polynomial time.*

Proof. Apply Lemma 4.15 to the operation f_1 for dimension $i = 2$. Then, there is an operation $f_1' \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_2(f_1')$ is canonical over $\mathrm{Aut}(\mathbb{Q}; <)$ and $m := \theta_1(f_1')$ equals min_3, mx_3, or mi_3. By Lemma 4.16, f_1' preserves $\neq_2$. Since $\theta_2(f_2)$ is an ll-operation, f_2 preserves $\neq_2$ as well. Therefore we may assume without loss of generality that $\mathfrak{D}$ contains the relation $\neq_2$. Since $\theta_1(\mathrm{Pol}(\mathfrak{D}))$ is closed (by Proposition 3.5), it follows from Proposition 4.6 that $\theta_1(\mathrm{Pol}(\mathfrak{D}))$ contains a pp-operation.

Let τ be the signature of $\mathfrak{D}$ and let $\mathfrak{A}$ be an instance of $\mathrm{CSP}(\mathfrak{D})$. For every $R \in \tau$ of arity k and $\bar{a} = (a_1, \ldots, a_k) \in R^{\mathfrak{A}}$, let $\phi_{R,\bar{a}}$ be the first-order definition of R in $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$, using the elements $a_1, \ldots, a_k$ as free variables. Since $\theta_1(\mathrm{Pol}(\mathfrak{D}))$ contains a pp-operation and $\theta_2(\mathrm{Pol}(\mathfrak{D}))$ contains an ll-operation, we may assume that $\phi_{R,\bar{a}}$ has the form described in Proposition 4.30. Let $\Phi = \{\phi_{R,\bar{a}} \mid R \in \tau \text{ and } \bar{a} \in R^{\mathfrak{A}}\}$. The set Φ can be computed in polynomial time since $\mathfrak{D}$ has a finite signature. It is clear that $\mathfrak{A}$ is a yes-instance of $\mathrm{CSP}(\mathfrak{D})$ if and only if Φ is satisfiable. We will now present a polynomial-time algorithm for checking the satisfiability of Φ; it is outlined in Algorithm 4.1. The basic idea is to compute two sets Ψ_1 and Ψ_2 of logical formulas that are simultaneously satisfiable if and only if Φ is satisfiable. The sets Ψ_1 and Ψ_2 are, in a sense that will be clarified below, connected to formulas in Φ that contain weakly 1-determined and 2-determined clauses, respectively.

Every $\phi \in \Phi$ is of the form $\phi_1 \wedge \psi_2$, where ϕ_1 is a conjunction of weakly 1-determined clauses and ψ_2 is a conjunction of clauses of the form

$$x_1 \neq_2 y_1 \vee \cdots \vee x_m \neq_2 y_m \vee z_1 <_2 z_0 \vee \cdots \vee z_\ell <_2 z_0 \vee (z_0 =_2 z_1 =_2 \cdots =_2 z_\ell).$$

Therefore $\hat{\psi}_2$ is a conjunction of ll-Horn clauses and by Theorem 4.12 preserved by an ll-operation. We let Ψ_2 denote the set of all formulas $\hat{\psi}_2$ obtained in this way from the members of Φ. Note that Ψ_2 can be computed in polynomial time since $\mathfrak{D}$ has a finite relational signature.

Since Ψ_2 is preserved by an ll-operation, we can check its satisfiability in polynomial time by Theorem 4.12 combined with Theorem 4.3. If Ψ_2 is not satisfiable, then we reject the input $\mathfrak{A}$.

Otherwise, we let E' denote the equality set of the instance. For each $(x, y) \in E'$, we put a pair (x^2, y^2) in the set E; this set will later on be used as an argument to the clause modification operator $\mathrm{cm}(\cdot)$ that was introduced in connection with Proposition 4.34.

Every ϕ in Φ is equivalent to a formula $\phi_1 \wedge \psi_2$ as described above. Note that the formula $\phi_1 \wedge \psi_2$ admits application of the clause modification operator $\mathrm{cm}(\cdot)$. Hence, we define Φ' to be the set of the formulas $\mathrm{cm}(\phi_1 \wedge \psi_2, E)$ obtained from formulas $\phi \in \Phi$. Note that each formula in Φ' still defines a relation that has a primitive positive definition in $\mathfrak{D}$ (since $\mathfrak{D}$ contains the relations $=_2$ and $\neq_2$). Further note that, up to renaming variables, only finitely many different formulas may appear in Φ': there are only finitely many inequivalent ways to remove $\neq_2$-literals or clauses with such literals and to add $=_2$- or $\neq_2$-conjuncts to the formulas $\phi_1 \wedge \psi_2$ defining one of the finitely many relations in $\mathfrak{D}$.

Every weakly 1-determined clause of a formula in Φ' that does not contain a literal of the form $x \neq_2 y$ is 1-determined. Every $\phi' \in \Phi'$ can thus be written as $\phi'_1 \wedge \phi'_2$, where ϕ'_i is a conjunction of i-determined clauses. It follows that ϕ' is equivalent to $\psi'_1 \wedge \phi'_2$ where $\psi'_1 = \mathrm{cr}(\phi', \{1\}, \phi'_1)$. Note that ψ'_1 is preserved by $\mathrm{Pol}(\mathfrak{D})$. Since $\mathfrak{D}$ has finite relational signature, there are only finitely many inequivalent formulas that can arise in this way so the formulas ψ'_1 can be computed in polynomial time: they can simply be stored in a fixed-size database that is computed off-line. By Lemma 3.11, $\hat{\psi}'_1$ is preserved by the operation $m \in \theta_1(\mathrm{Pol}(\mathfrak{D}))$. Let Ψ_1 be the set of all formulas $\hat{\psi}'_1$ obtained from Φ' in this way. We may use the algorithmic part of Theorem 4.3 to decide whether Ψ_1 is satisfiable. If Ψ_1 is not satisfiable, then we reject the input $\mathfrak{A}$ and we accept it otherwise. We claim that in this case Φ is satisfiable and, consequently, that $\mathfrak{A}$ has a homomorphism to $\mathfrak{D}$.

Indeed, let $s\colon A \to \mathbb{Q}$ be a solution to Ψ_1 and let $t\colon A \to \mathbb{Q}$ be a solution to Ψ_2. Since $\mathfrak{D}$ is preserved by an operation g such that $\theta_2(g)$ is an ll-operation, and ll-operations are injective, we may assume that t satisfies $x \neq_2 y$ unless this literal has been removed from Φ by the algorithm. Then the map $x \mapsto (s(x), t(x))$ satisfies all formulas in Φ and it follows that $\mathfrak{A}$ admits a homomorphism to $\mathfrak{D}$. ◻

4.4.2. Classification Result. The known results about first-order expansions of $(\mathbb{Q}; <)$ from Section 4.1 combined with the results from Section 4.3 imply an algebraic dichotomy for polymorphism clones of first-order expansions of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$; the algebraic dichotomy implies a complexity dichotomy, using the results from Section 4.4.1.

THEOREM 4.35. *Let $\mathfrak{D}$ be a first-order expansion of $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$. Exactly one of the following two cases applies.*

- *For each $i \in \{1, 2\}$ we have that $\theta_i(\mathrm{Pol}(\mathfrak{D}))$ contains min_3, mx_3, mi_3, or ll_3, or one of their duals. Furthermore, $\mathfrak{D}$ has a pwnu polymorphism and if $\mathfrak{D}$ has a finite relational signature, then $\mathrm{CSP}(\mathfrak{D})$ is in P.*
- *$\mathrm{Pol}(\mathfrak{D})$ has a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$. In this case, $\mathfrak{D}$ has a finite-signature reduct whose CSP is NP-complete.*

Proof. For $i \in \{1, 2\}$, let $\mathscr{C}_i = \theta_i(\mathrm{Pol}(\mathfrak{D}))$. If $\mathscr{C}_i$, for some $i \in \{1, 2\}$, has a uniformly continuous minor-preserving map from $\mathscr{C}_i$ to $\mathrm{Pol}(K_3)$, then by composing uniformly continuous minor-preserving maps there is also a uniformly continuous minor-preserving map from $\mathrm{Pol}(\mathfrak{D})$ to $\mathrm{Pol}(K_3)$, which implies that $\mathfrak{D}$ has a finite signature reduct whose CSP is NP-hard by Corollary 2.12. Assume henceforth that there is no uniformly continuous minor-preserving map from $\mathscr{C}_i$ to $\mathrm{Pol}(K_3)$.

Let $\mathfrak{D}_i$ be the structure with domain $\mathbb{Q}$ that contains all relations that are preserved by $\mathscr{C}_i$; note that $\mathrm{Pol}(\mathfrak{D}_i) = \mathscr{C}_i$ by Proposition 3.5. Clearly, $\mathscr{C}_i$ contains $\mathrm{Aut}(\mathbb{Q}; <)$ and preserves $<$, so $\mathfrak{D}_i$ is a first-order expansion of $(\mathbb{Q}; <)$. By Theorem 4.3, $\mathscr{C}_i$ contains min_3, mx_3, mi_3, or ll_3, or the

Algorithm 4.1 Solve-by-Factors

Input: an instance $\mathfrak{A}$ of CSP($\mathfrak{D}$)
$\Phi \coloneqq \{\phi_{R,\bar{a}} \mid R \in \tau, \bar{a} \in R^{\mathfrak{A}}\}$
for all $\phi \in \Phi$ **do**
  write ϕ as $\phi_1 \wedge \psi_2$, where ϕ_1 is a conjunction of weakly 1-determined clauses and ψ_2 is a conjunction of 2-determined clauses such that each clause of $\hat{\psi}_2$ is ll-Horn
end for
$\Psi_2 \coloneqq \{\hat{\psi}_2 \mid \phi \in \Phi\}$
▷ *Satisfiability of* Ψ_2 *can be checked in polynomial time by Theorems 4.3 and 4.12.*
if Ψ_2 is not satisfiable **then**
  reject
end if
let E' be the equality set of Ψ_2
$E \coloneqq \{(x^2, y^2) \mid (x, y) \in E'\}$
for all $\phi \in \Phi$ **do**
  write $\mathrm{cm}(\phi, E)$ as $\bigwedge_{p \in \{1,\dots,n\} \setminus S} \phi_1' \wedge \psi_2 \wedge \psi_2'$, where ϕ_1' is the conjunction of 1-determined clauses resulting from ϕ_1 and ψ_2' is a conjunction of the added conjuncts
end for
$\Phi' \coloneqq \{\mathrm{cm}(\phi, E) \mid \phi \in \Phi\}$
▷ *Every* $\phi' \in \Phi'$ *defines a relation that is primitively positively definable over* $\mathfrak{D}$.
$\Psi_1 \coloneqq \{\hat{\psi}_1 \mid \phi' \in \Phi', \psi_1 = \mathrm{cr}(\phi', \{1\}, \phi_1')\}$
▷ Ψ_1 *is preserved by* min_3, mx_3, *or* mi_3 *and its satisfiability can be checked in polynomial time by Theorems 4.2 and 4.3.*
if Ψ_1 is not satisfiable **then**
  reject
end if
accept

dual of one of these operations. By Lemma 4.20, we may assume that $\mathscr{C}_i$ contains an operation $f_i \in \{\mathrm{min}_3, \mathrm{mx}_3, \mathrm{mi}_3, \mathrm{ll}_3\}$; we assume without loss of generality that $f_i = \mathrm{ll}_3$ whenever $\mathscr{C}_i$ contains ll_3. If $\mathscr{C}_i$, for some $i \in \{1, 2\}$, contains a lex-operation, then $\mathscr{C}_i$ contains ll_3: otherwise, $\theta_i(\mathrm{Pol}(\mathfrak{D}))$ would have to contain mi_3, mx_3, or mi_3, and by Proposition 4.6 a pp-operation. It then follows from the final statement in Theorem 4.12 that $\mathscr{C}_i$ contains ll_3 and this contradicts our assumptions.

Assume that $\mathscr{C}_2$ contains ll_3, and hence an ll-operation by Theorem 4.12. If $\mathscr{C}_1$ contains min_3, mi_3, or mx_3, then the polynomial-time tractability of CSP($\mathfrak{D}$) follows from Proposition 4.34. Otherwise, $\mathscr{C}_1$ contains ll_3 (and hence an ll-operation by Theorem 4.12) and the polynomial-time tractability of CSP($\mathfrak{D}$) follows from Proposition 4.33. The case when $\mathscr{C}_1$ contains ll_3 follows from the same argument with the roles of the two dimensions exchanged.

Suppose in the following that neither $\mathscr{C}_1$ nor $\mathscr{C}_2$ contains a lex-operation. Remark 4.11 combined with Theorem 4.12 imply that they do not contain the operation ll_3. Then, both $\mathscr{C}_1$ and $\mathscr{C}_2$ contain min_3, mi_3, or mx_3, and Proposition 4.6 imply that both $\mathscr{C}_1$ and $\mathscr{C}_2$ contain a pp-operation. By Proposition 4.31, every relation of $\mathfrak{D}$ can be defined by a conjunction of 1-determined clauses and of 2-determined clauses. Thus, Lemma 3.12 implies that $\mathrm{Pol}(\mathfrak{D})$ contains $\mathscr{C}_1 \times \mathscr{C}_2$. Then the polynomial-time tractability of CSP($\mathfrak{D}$) follows from Corollary 3.8 applied to $\mathfrak{D}_1, \mathfrak{D}_2$, and $\mathfrak{D}$.

We continue by proving that $\mathfrak{D}$ admits a pwnu polymorphism. Let f be the ternary operation such that $\theta_i(f)$ equals f_i for every $i \in \{1,2\}$. We claim that f preserves $\mathfrak{D}$. Let ϕ be a formula that defines a relation from $\mathfrak{D}$. By Proposition 4.26, we may assume that ϕ is a normal conjunction of clauses each of which is weakly i-determined for some i. Let a,b,c be tuples that satisfy ϕ. Let ψ be a clause of ϕ. We may assume that ψ is weakly 2-determined, since the case where it is weakly 1-determined can be treated analogously. Then ψ is of the form $\psi' \vee \psi''$, where ψ' is a 2-determined clause and ψ'' is a disjunction of $\neq_1$-literals. We show that $f(a,b,c)$ satisfies ψ.

Note that it follows from the discussion above that $\mathscr{C}_i$ contains for each $i \in \{1,2\}$ a pp-operation or an ll-operation. If ψ contains a literal $x \neq_1 y$, then it is not 2-determined and it follows from Proposition 4.31 that either $\mathscr{C}_1$ does not contain a pp-operation or $\mathscr{C}_2$ contains a lex-operation. In the first case, $\mathscr{C}_1$ does not contain min_3, mx_3 or mi_3 by Proposition 4.6 so $f_1 = \text{ll}_3$. In the second case, Proposition 4.27 implies that ψ is weakly 1-determined. We see that $f_2 = \text{ll}_3$ by Theorem 4.12 since either (1) $\mathscr{C}_2$ simultaneously contains a lex-operation and a pp-operation or (2) $\mathscr{C}_2$ contains an ll-operation. We may therefore assume that $f_1 = \text{ll}_3$ since, otherwise, we can treat ψ as a weakly 1-determined clause.

If one of a,b,c satisfies the literal $x \neq_1 y$, then $f(a,b,c)$ satisfies the literal as well since $\theta_1(f) = f_1$ is injective. So suppose that none of a,b,c satisfies such literals. We show that $f(a,b,c)$ satisfies ψ'. Since $f_2 \in \mathscr{C}_2$, there is $f_2' \in \text{Pol}(\mathfrak{D})$ such that $\theta_2(f_2') = f_2$. Since f_2' preserves ϕ and the relation $=_1$, the tuple $f_2'(a,b,c)$ must satisfy the 2-determined clause ψ' and hence $f_2(a,b,c)$ satisfies $\hat{\psi}'$ by Lemma 3.11. Another application of Lemma 3.11 shows that $f(a,b,c)$ satisfies ψ' as well.

Finally, we prove that f is indeed a pwnu polymorphism of $\mathfrak{D}$. If e_1^i, e_2^i, e_3^i show that f_i is a pwnu polymorphism of $\mathfrak{D}_i$, then $e_j := (e_j^1, e_j^2)$, for $j \in \{1,2,3\}$, are the endomorphisms of $\mathfrak{D}$ that show that f is a pwnu polymorphism of $\mathfrak{D}$.

Since $\overline{\text{Aut}(\mathbb{Q};<)} = \text{End}(\mathbb{Q};<)$, it follows from Lemma 3.9 that $\overline{\text{Aut}(\mathfrak{D})} = \text{End}(\mathfrak{D})$. Therefore, Lemma 2.13 implies that the two cases in the statement are mutually exclusive. ∎

4.5. Classification in the n-Dimensional Case. The approach in the previous section can be generalised to first-order expansions of $(\mathbb{Q};<)^{(n)} = (\mathbb{Q}^n; <_1, =_1, \ldots, <_n, =_n)$. However, this requires some work both on the algebraic and the algorithmic side. In order to make the transition as smooth as possible, we generalize weakly i-determined clauses into *S-weakly i-determined* clauses. This allows us to systematically transfer both our 2-dimensional algebraic results (Propositions 4.27-4.32) and algorithmic results (Propositions 4.33 and 4.34 into the n-dimensional setting. From now on and for the remainder of Section 4, we let the symbol $\mathfrak{D}$ denote a first-order expansions of $(\mathbb{Q};<)^{(n)} = (\mathbb{Q}^n; <_1, =_1, \ldots, <_n, =_n)$. We begin by generalising Definition 4.25.

Definition 4.36. *Let $S \subseteq \{1,\ldots,n\}$ and $p \in \{1,\ldots,n\}$. A clause is called* S-weakly p-determined *if it is of the form*

$$\psi \vee \bigvee_{i \in \{1,\ldots,k\} \text{ and } j_i \in S} x_i \neq_{j_i} y_i$$

where ψ is p-determined and $k \geq 0$. A clause is called weakly p-determined *if it is $\{1,\ldots,n\} \setminus \{p\}$-weakly p-determined (note that this is consistent with the notion of weakly p-determined for $n=2$ from Definition 4.25).*

Next, we connect conjunctions of S-weakly p-determined clauses with first-order expansions of $(\mathbb{Q};<)^{(n)}$ that admit certain polymorphisms. Recall the normal forms of formulas over $(\mathbb{Q};<)^{(n)}$ defined at the end of Section 4.3.1; we will use them in similar fashion as in the case $n=2$.

Proposition 4.37. *Suppose that for every $p \in \{1,\dots,n\}$ there is an operation $f \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(f)$ is an* ll*-operation or a* pp*-operation. Then for every relation R of $\mathfrak{D}$, if ϕ is a first-order definition of R over $(\mathbb{Q}^n; <_1, =_1, \dots, <_n, =_n)$ that is normal, then ϕ is a conjunction of clauses each of which is weakly i-determined for some $i \in \{1,\dots,n\}$.*

Proof. The proof is a generalisation of the proof of Proposition 4.26; the key step is to use the generalisation of Corollary 4.19 which states that for $i, j \in \{1,\dots,n\}$ and $f \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_i(f)$ is an ll- or a pp-operation there is an operation $g \in \mathrm{Pol}(\mathfrak{D}; \leq_1, \neq_1, \dots, \leq_n, \neq_n)$ such that

- $\theta_i(g)$ is dominated by the first argument (or even equal to π^2_1 if $\theta_i(f)$ is a pp-operation) and
- $\theta_j(g)$ is dominated by the second argument.

To prove this, note first that we may without loss of generality assume that f preserves $\leq_k$ and $\neq_k$ for all k; for $k = i$ this follows from the assumption and for $k \neq i$ we may repeatedly apply the n-dimensional generalisation of Lemma 4.17 (as discussed immediately after Lemma 4.15) to obtain an operation that preserves $\leq_k$ and $\neq_k$. By applying Lemma 4.17 to canonise the operation in the j-th position and subsequent application of the n-dimensional generalisation of Lemma 4.18 to modify the i-th position, we can prove the statement analogously to the proof of Corollary 4.19.

To prove the proposition, one can proceed as in the proof of Proposition 4.26; the only difference is that the choice of the polymorphism g depends on the considered pair of literals, because it needs to have the domination property in the right dimensions. In fact, the generalisation of that proof yields the conclusion under the weaker assumption that for all but at most one $p \in \{1,\dots,n\}$ there exists $f \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(f)$ is an ll-operation, a pp-operation, or the dual of such an operation. ☐

The following proposition gives a concrete syntactic description for defining formulas of relations of $\mathfrak{D}$ based on the polymorphisms of $\mathfrak{D}$. It can be viewed as a generalisation of Propositions 4.27-4.32.

Proposition 4.38. *Let $S \subseteq \{1,\dots,n\}$ be such that*

- *for every $p \in S$ there exists $f_p \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(f_p)$ is an* ll*-operation, and*
- *for every $p \in \{1,\dots,n\} \setminus S$ there exists $f_p \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(f_p)$ is a* pp*-operation, but there is no $g \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(g)$ is a* lex*-operation.*

Then, the following hold:

1. *every relation of $\mathfrak{D}$ can be defined by a conjunction of clauses each of which is an S-weakly p-determined clause for some $p \in \{1,\dots,n\}$,*
2. *if $p \in \{1,\dots,n\} \setminus S$, then the S-weakly p-determined clauses can be chosen to be of the form*

$$u_1 \neq_{i_1} v_1 \vee \cdots u_m \neq_{i_m} v_m \vee y_1 \neq_p x \vee \cdots \vee y_k \neq_p x \vee z_1 \leq_p x \vee \cdots \vee z_l \leq_p x, \tag{4.5}$$

   where $i_1, \dots, i_m \in S$, and
3. *if $p \in S$, then the S-weakly p-determined clauses can be chosen to be* ll*-Horn clauses of the form*

$$x_1 \neq_{i_1} y_1 \vee \cdots \vee x_m \neq_{i_m} y_m \vee z_1 <_p z_0 \vee \cdots \vee z_\ell <_p z_0 \vee (z_0 =_p z_1 =_p \cdots =_p z_\ell)$$

   for $i_1, \dots, i_m \in S$ (and where the last disjunct may not appear).

Proof. Every relation of $\mathfrak{D}$ has a definition by a normal formula and, by Proposition 4.37, it can be defined by a conjunction of clauses each of which is weakly i-determined for some $i \in \{1,\dots,n\}$.

Let R be a relation of $\mathfrak{D}$ and let ϕ be such a definition. We show step by step that the statements in items $1-3$ hold true for R.

Proof of item 1. Let ψ be a weakly i-determined clause of ϕ for some $i \in \{1,\dots,n\}$. When $i \in S$ and $j \in \{1,\dots,n\} \setminus S$, then we can proceed as in the proof of Proposition 4.27: we use the generalised version of Corollary 4.19 described in detail in the previous proof and we rule out the possibility that ψ simultaneously contains a $\{<_i, =_i\}$-literal and a $\neq_j$-literal. Therefore, ψ is an S-weakly i-determined clause or an S-weakly j-determined clause in this case.

If $i, j \in \{1,\dots,n\} \setminus S$ are distinct, then there is an operation $g \in \mathrm{Pol}(\mathfrak{D}; \leq_1, \neq_1, \dots, \leq_n, \neq_n)$ with $\theta_i(g) = \pi_1^2$ and $\theta_j(g) = \pi_2^2$. To see this, we first assume that g preserves $\leq_k$ and $\neq_k$ for all k — if this is not the case, then we repeatedly apply Lemma 4.17 in all but the i-th dimension. The existence of g can now be proved similarly as in the proof of Proposition 4.31 (the proof uses the fact that there is no $g \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_i(g)$ or $\theta_j(g)$ is a lex-operation). As in the proof of Lemma 3.12, g can be used to prove that ψ cannot contain a $\{<_i, =_i, \neq_i\}$-literal and $\neq_j$-literal at the same time. Hence, the clause ψ is an S-weakly i-determined clause. It follows that the normal formula ϕ is in fact a conjunction of clauses each of which is S-weakly p-determined for some p.

Proof of item 2. We now prove that we may choose the clauses that are S-weakly p-determined for $p \notin S$ to have the syntactic form (4.5). We will proceed analogously to the proof of Proposition 4.29. Let $p \in \{1,\dots,n\} \setminus S$ and $\phi = \phi_p \wedge \phi_0$, where ϕ_p is the conjunction of all S-weakly p-determined clauses of ϕ and ϕ_0 is the conjunction of the remaining clauses. Let ϕ'_p be the conjunction of all clauses of the form (4.5) that are reduced and implied by ϕ. We will show that $\phi' = \phi'_p \wedge \phi_0$ implies ϕ and hence defines R. Applying the same procedure for all $p \in \{1,\dots,n\} \setminus S$ concludes the proof of item 2.

To see that ϕ' implies ϕ, we use the same orbit argument as in the proof of Proposition 4.29: ϕ_p is equivalent to a conjunction of S-weakly p-determined clauses of the form

$$\chi \vee y_1 \circ_1 y_2 \vee \cdots \vee y_k \circ_k y_{k+1},$$

where $\chi = \bigvee_{j=1}^{m} u_j \neq_{i_j} v_j$, $i_j \in S$, $j = 1,\dots,m$, and $\circ_1,\dots,\circ_k \in \{\neq_p, \geq_p\}$. We assume that these clauses are minimal in the same sense as in the proof of Proposition 4.29. Let ψ be such a clause of ϕ. The argument in the rest of the proof of Proposition 4.29 is not dependent on the indices i_j in the literals $u_j \neq_{i_j} v_j$ in χ. Thus, it is applicable also in this case and shows that ψ is implied by ϕ'. This proves that ϕ' implies ϕ since ψ was chosen arbitrarily.

Proof of item 3. By items 1 and 2, we may assume without loss of generality that ϕ satisfies the following condition: for every $p \in \{1,\dots,n\} \setminus S$, every S-weakly p-determined clause of ϕ is of the form (4.5). Let ϕ_1 be a conjunction of all S-weakly p-determined clauses of ϕ where $p \in S$ and let ϕ_2 be the conjuction of the remaining clauses of ϕ. We will now use Corollary 3.16 and the conjunction replacement operator $\mathrm{cr}(\cdot)$. We note that ϕ_1 is a conjunction of S-determined clauses and hence ϕ is equivalent to a formula $\psi_1 \wedge \phi_2$ where $\psi_1 = \mathrm{cr}(\phi, S, \phi_1)$. Without loss of generality, we may assume that ψ_1 is normal. By Proposition 4.37, ψ_1 is in fact a conjunction of clauses each of which is weakly p-determined for some p and hence S-weakly p-determined for some $p \in S$.

Recall the variable expansion operator $\mathrm{ve}(\cdot)$ that we defined just before Proposition 4.32. Since for every $p \in S$, there is an operation $f_p \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(f_p)$ is an ll-operation and ψ_1 is S-determined, it can be shown analogously to the claim in the proof of Proposition 4.32 that $\mathrm{ve}(\psi_1)$ is preserved by every ll-operation. By Theorem 4.12, $\mathrm{ve}(\psi_1)$ is equivalent to a conjunction of ll-Horn clauses. Since the formula $\psi_1 \wedge \phi_2$ defines R, item 3 follows. ∎

Note that the formula produced by Proposition 4.38 is not necessarily normal. Also note the difference between the proof for $n = 2$ and general n: For $n = 2$, there are just three cases – S is empty, $S = \{p\}$ for some p, or $S = \{1, 2\}$. If $S = \{p\}$, then S-determined clauses are p-determined. If $S = \{1, 2\}$, then the formula ϕ_1 is equal to ϕ and thus trivially preserved by $\mathrm{Pol}(\mathfrak{D})$.

We continue by generalizing our algorithmic results. Recall the clause modification operator cm from Section 4.4.1. We generalize this operator for the case of arbitrary $n > 2$. Let X be a set of variables, $S \subseteq \{1, \ldots n\}$ and $E \subseteq \{(x^i, y^i) \mid x, y \in X, i \in S\}$ where x^i, y^i denote fresh variables. Assume ϕ is a formula over $(\mathbb{Q}^n, <_1, =_1, \ldots, <_n, =_n)$ that only uses variables from X. If ϕ equals $\bigwedge_{p\in\{1,\ldots,n\}\setminus S} \phi_p \wedge \phi_S$, where ϕ_p is a conjunction of S-weakly p-determined clauses for each p and ϕ_S is a conjunction of S-determined clauses, then $\mathrm{cm}(\phi, E)$ denotes the formula resulting from ϕ by performing the following procedure for each $p \in \{1, \ldots, n\} \setminus S$:

- If $(x^i, y^i) \in E$ and ϕ_p contains a clause with the literal $x \neq_i y$, then remove this literal from the clause and add the conjunct $x =_i y$ to ϕ.
- If $(x^i, y^i) \notin E$ and ϕ_p contains a clause with the literal $x \neq_i y$, then delete all literals in the clause but this one.

Note that the previously defined notion of the operator cm was a special case for $n = 2$ and $S = \{i\} \subseteq \{1, 2\}$.

The proof of our n-dimensional computational result is based on the ideas underlying the 2-dimensional results found in Propositions 4.33 and 4.34, but there are also crucial differences. It may be enlightening to compare the generalized Algorithm 4.2 with the two-dimensional Algorithm 4.1 underlying Proposition 4.34. In particular, such a comparison may clarify how S-weakly p-determined and the n-dimensional cm operator come into play.

PROPOSITION 4.39. *Suppose that $\mathfrak{D}$ has a finite relational signature τ and suppose that for each $p \in \{1, \ldots, n\}$ there exists $f_p \in \mathrm{Pol}(\mathfrak{D})$ such that $\theta_p(f_p)$ equals* ll_3*,* min_3*,* mx_3*, or* mi_3*. Then,* $\mathrm{CSP}(\mathfrak{D})$ *can be solved in polynomial time.*

Proof. Without loss of generality, suppose that for every $p \in \{1, \ldots, n\}$ such that $\theta_p(\mathrm{Pol}(\mathfrak{D}))$ contains ll_3, the operation f_p is chosen to be such that $\theta_p(f_p) = \mathrm{ll}_3$. Let $S \subseteq \{1, \ldots, n\}$ be the set of all such p. Moreover, we may assume that $\mathfrak{D}$ contains relations $\neq_i$ for every $i \in S$. Otherwise we repeatedly apply Lemma 4.15 on the operations f_p, $p \in \{1, \ldots, n\}$, and obtain polymorphisms f'_p such that

- $\theta_i(f'_p)$ is canonical over $\mathrm{Aut}(\mathbb{Q}; <)$ for every $i \neq p$ (and hence preserves $\neq_i$ by Lemma 4.16),
- $\theta_p(f_p)$ equals ll_3, min_3, mx, or mi_3 (and hence preserves $\neq_p$ whenever $p \in S$).

Note that $\theta_p(\mathrm{Pol}(\mathfrak{D}))$ contains a pp-operation for every $p \in \{1, \ldots, n\} \setminus S$ (by Proposition 3.5 and Proposition 4.6), but it does not contain a lex-operation (by Theorem 4.12). By Theorem 4.12, $\theta_p(\mathrm{Pol}(\mathfrak{D}))$ contains a ll-operation for $p \in S$. Let $\mathfrak{A}$ be an instance of $\mathrm{CSP}(\mathfrak{D})$. For every $R \in \tau$ of arity k and $\bar{a} = (a_1, \ldots, a_k) \in R^{\mathfrak{A}}$, let $\phi_{R,\bar{a}}$ be the first-order definition of R in the structure $(\mathbb{Q}^n; <_1, =_1, \ldots, <_n, =_n)$ using the elements $a_1, \ldots, a_k$ as the free variables. We then may assume that $\phi_{R,\bar{a}}$ is of the form as described in Proposition 4.38.

In Algorithm 4.2 we present the algorithm for deciding whether a given instance $\mathfrak{A}$ of $\mathrm{CSP}(\mathfrak{D})$ has a homomorphism to $\mathfrak{D}$. The algorithm uses the same ideas as are used in the proof of Proposition 4.34. In particular, we can show by similar arguments that the sets Φ, Ψ_S, Φ' and Ψ_p, $p \in \{1, \ldots, n\} \setminus S$, can be computed in polynomial time. Finally, the set E can be computed in polynomial time (as was pointed out after Proposition 4.33) so the whole algorithm runs in polynomial time.

It is clear that if the algorithm rejects, then there is no homomorphism from $\mathfrak{A}$ to $\mathfrak{D}$. In case that the algorithm accepts, the existence of a homomorphism from $\mathfrak{A}$ to $\mathfrak{D}$ can be proved in a

Algorithm 4.2 Solve-by-Factors-n

Input: an instance $\mathfrak{A}$ of CSP($\mathfrak{D}$)
$\Phi \coloneqq \{\phi_{R,\bar{a}} \mid R \in \tau, \bar{a} \in R^{\mathfrak{A}}\}$
for all $\phi \in \Phi$ **do**
  write ϕ as $\bigwedge_{p\in\{1,\dots,n\}\setminus S} \phi_p \wedge \phi_S$, where ϕ_p is a conjunction of S-weakly p-determined clauses and ϕ_S is a conjunction of S-determined ll-Horn clauses
end for
$\Psi_S \coloneqq \{\mathrm{ve}(\phi_S) \mid \phi \in \Phi\}$
▷ *Ψ_S contains conjunctions of* ll*-Horn clauses over* $(\mathbb{Q}, <)$ *and its satisfiability can be checked in polynomial time by Theorems 4.3 and 4.12.*
if Ψ_S is not satisfiable **then**
  reject
end if
let E denote the equality set corresponding to Ψ_S
for all $\phi \in \Phi$ **do**
  write $\mathrm{cm}(\phi, E)$ as $\bigwedge_{p\in\{1,\dots,n\}\setminus S} \phi'_p \wedge \phi_S \wedge \phi'_S$, where ϕ'_p is the conjunction of p-determined clauses resulting from ϕ_p and ϕ'_S is a conjunction of the added conjuncts
end for
$\Phi' \coloneqq \{\mathrm{cm}(\phi, E) \mid \phi \in \Phi\}$
for all $p \in \{1, \dots, n\} \setminus S$ **do**
  ▷ *Every $\phi' \in \Phi'$ defines a relation that is primitively positively definable over $\mathfrak{D}$.*
  $\Psi_p \coloneqq \{\hat{\psi}_p \mid \phi' \in \Phi', \psi_p = \mathrm{cr}(\phi', \{p\}, \phi'_p)\}$
  ▷ *Ψ_p is preserved by* min_3, mx_3, *or* mi_3.
  ▷ *Satisfiability of Ψ_p can be checked in polynomial time by Theorems 4.2 and 4.3.*
  if Ψ_p is not satisfiable **then**
    reject
  end if
end for
accept

similar fashion as in the proof of Proposition 4.34. □

We are now in the position of proving the main result of this section.

THEOREM 4.40. *Let $\mathfrak{D}$ be a first-order expansion of $(\mathbb{Q}; <_1, =_1, \dots, <_n, =_n)$. Exactly one of the following two cases applies.*

- *For each $p \in \{1, \dots, n\}$ we have that $\theta_p(\mathrm{Pol}(\mathfrak{D}))$ contains* min_3, mx_3, mi_3, *or* ll_3, *or one of their duals. In this case, $\mathfrak{D}$ has a pwnu polymorphism. If $\mathfrak{D}$ has a finite relational signature, then* CSP($\mathfrak{D}$) *is in P.*
- Pol($\mathfrak{D}$) *has a uniformly continuous minor-preserving map to* Pol(K_3). *In this case $\mathfrak{D}$ has a finite-signature reduct whose* CSP *is NP-complete.*

Proof. For $p \in \{1, \dots, n\}$, define $\mathscr{C}_p \coloneqq \theta_p(\mathrm{Pol}(\mathfrak{D}))$. If $\mathscr{C}_p$, for some $p \in \{1, \dots, n\}$, has a uniformly continuous minor-preserving map from $\mathscr{C}_p$ to Pol(K_3), then by composing uniformly continuous minor-preserving maps there is also a uniformly continuous minor-preserving map from Pol($\mathfrak{D}$) to Pol(K_3). This implies that $\mathfrak{D}$ has a finite-signature reduct whose CSP is NP-hard by Corollary 2.12.

Otherwise, for every $p \in \{1,\dots,n\}$ the clone $\mathscr{C}_p$ does not have a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$. Since $\mathscr{C}_p$ is a closed clone (by Proposition 3.5) that contains $\mathrm{Aut}(\mathbb{Q};<)$ and preserves $<$, there exists a first-order expansion $\mathfrak{D}_p$ of $\mathrm{Aut}(\mathbb{Q};<)$ such that $\mathrm{Pol}(\mathfrak{D}_p) = \mathscr{C}_p$. We may therefore apply Theorem 4.3 and conclude that for every $p \in \{1,\dots,n\}$ the clone $\mathscr{C}_p$ contains an operation f_p which equals min_3, mx_3, mi_3, or ll_3, or the dual of one of these operations. By the version of Lemma 4.20 for n-fold algebraic products, we may assume without loss of generality that $f_p \in \{\mathrm{min}_3, \mathrm{mx}_3, \mathrm{mi}_3, \mathrm{ll}_3\}$ for every $p \in \{1,\dots,n\}$. In case that $\mathfrak{D}$ has a finite relational signature, the polynomial-time tractability of $\mathrm{CSP}(\mathfrak{D})$ follows from Proposition 4.39.

We may assume that $f_p = \mathrm{ll}_3$ for every p such that $\mathscr{C}_p$ contains ll_3. Let f be the ternary operation such that $\theta_p(f)$ equals f_p for every $p \in \{1,\dots,n\}$. We claim that f preserves $\mathfrak{D}$. Let ϕ be a formula that defines a relation from $\mathfrak{D}$ and has the form as described in Proposition 4.38 (the argument in the proof of Proposition 4.39 implies that the assumptions are satisfied). We show that f preserves ϕ. Let a, b, c be tuples that satisfy ϕ and let ψ be a clause of ϕ. We see that there is a $p \in \{1,\dots,n\}$ such that ψ is S-weakly p-determined. As in the proof of Theorem 4.35, one can show that $f(a,b,c)$ satisfies ψ. It follows that f preserves ϕ and $f \in \mathrm{Pol}(\mathfrak{D})$.

As in the case when $n = 2$ (Theorem 4.35), we can show that f is a pwnu polymorphism, and hence it follows from Lemma 3.9 and Lemma 2.13 that the two cases of the statement are mutually exclusive. ☐

4.6. Classification of Binary Relations. A relational signature is called *binary* if all its relation symbols have arity two, and a relational structure is binary if its signature is binary. If $\mathfrak{D}$ is binary, then the results from the previous sections can be substantially strengthened. Note that an ω-categorical structure has only finitely many distinct relations of arity at most two so we may assume that binary structures have a finite signature.

DEFINITION 4.41. *A formula is called an* Ord-Horn clause *if it is of the form*

$$x_1 \neq y_1 \vee \cdots \vee x_m \neq y_m \vee z_1 \circ z_0$$

where $\circ \in \{<, \leq, =\}$*, it is permitted that* $m = 0$*, and the final disjunct may be omitted. An* Ord-Horn formula *is a conjunction of Ord-Horn clauses.*

Ord-Horn clauses are ll-Horn formulas and a first-order formula over $(\mathbb{Q};<)$ is equivalent to an Ord-Horn formula if and only if it is preserved by an ll-operation and the dual of an ll-operation [32]. We say that a relation has an Ord-Horn definition if it can be defined by an Ord-Horn formula. The polynomial-time tractability of $\mathrm{CSP}(\mathfrak{B})$ if all relations of $\mathfrak{B}$ have an Ord-Horn definition follows from Theorem 4.3 and Theorem 4.12, but this was first shown by Nebel and Bürckert [87] using a very different approach. The reader should note that the Ord-Horn fragment does not have a characterisation in terms of equations satisfied by the polymorphism clone [39, Theorem 7.2].

The following theorem strengthens the results of Theorem 4.40 for binary first-order expansions, providing a very concrete syntactic condition for tractability based on Ord-Horn formulas.

THEOREM 4.42. *Let* $\mathfrak{D}$ *be a binary first-order expansion of* $(\mathbb{Q}; <_1, =_1, \dots, <_n, =_n)$*. Then exactly one of the following two cases applies.*

- *Each relation in* $\mathfrak{D}$*, viewed as a relation of arity* $2n$ *over* $\mathbb{Q}$*, has an Ord-Horn definition. In this case,* $\mathfrak{D}$ *has a pwnu polymorphism and* $\mathrm{CSP}(\mathfrak{D})$ *is in P.*
- $\mathrm{Pol}(\mathfrak{D})$ *has a uniformly continuous minor-preserving map to* $\mathrm{Pol}(K_3)$*. In this case,* $\mathrm{CSP}(\mathfrak{D})$ *is NP-complete.*

Proof. If the second item of the statement does not apply, then Theorem 4.40 implies that $\mathfrak{D}$ has a pwnu polymorphism and for each $p \in \{1,\dots,n\}$ we have that $\theta_p(\mathrm{Pol}(\mathfrak{D}))$ contains min_3, mx_3, mi_3, or ll_3, or one of their duals. By Lemma 4.20 we may focus on the situation that $\theta_p(\mathrm{Pol}(\mathfrak{D}))$ contains min_3, mx_3, mi_3, or ll_3 (note that the dual of a relation with an Ord-Horn definition has an Ord-Horn definition as well). Then Proposition 4.6 and Theorem 4.12 imply that the assumptions of Proposition 4.37 hold, and therefore every relation of $\mathfrak{D}$ can be defined by a normal conjunction of clauses each of which is weakly s-determined for some $s \in \{1,\dots,n\}$. If such a clause contains two disjuncts of the form $x <_i y$ and $y <_i x$, then replace the disjuncts by $x \neq_i y$. If such a clause contains two disjuncts of the form $x =_i y$ and $x \neq_i y$, then remove the clause (since it is always true). If such a clause contains two disjuncts of the form $x <_i y$ and $x =_i y$, then replace the disjuncts by $x \leq_i y$. Since the relations of $\mathfrak{D}$ are binary, the resulting formula is Ord-Horn. The result follows since we know that the satisfiability of Ord-Horn formulas can be decided in polynomial time. Lemma 2.13 implies that the two items cannot hold simultaneously. ∎

5. Complexity Classification Transfer. Assume that $\mathcal{C}$ and $\mathcal{D}$ are classes of structures and that the complexity of $\mathrm{CSP}(\mathfrak{D})$ is known for every $\mathfrak{D} \in \mathcal{D}$. A *complexity classification transfer* is a process that systematically uses this information for inferring the complexity of $\mathrm{CSP}(\mathfrak{C})$ for every $\mathfrak{C} \in \mathcal{C}$. The particular method that we will use originally appeared in [19]. Combined with our classification for first-order expansions of $(\mathbb{Q};<)^{(n)}$, this method allows us to derive several new dichotomy results in Section 6.

Let $\mathfrak{C}$ and $\mathfrak{D}$ denote relational structures. Recall that an interpretation of $\mathfrak{C}$ in $\mathfrak{D}$ is a partial surjection from a finite power of D to C. Two interpretations I and J of $\mathfrak{C}$ in $\mathfrak{D}$ are called *primitively positively homotopic*[1] (pp-homotopic) if the relation $\{(\bar{x},\bar{y}) \mid I(\bar{x}) = J(\bar{y})\}$ is primitively positively definable in $\mathfrak{D}$. The *identity interpretation* of a τ-structure $\mathfrak{C}$ is the identity map on C, which is clearly a primitive positive interpretation. We write $I_1 \circ I_2$ for the natural composition of two interpretations I_1 and I_2.

Definition 5.1. *Let $\mathfrak{C}$ and $\mathfrak{D}$ be two mutually primitively positively interpretable structures with a primitive positive interpretation I of $\mathfrak{C}$ in $\mathfrak{D}$ and a primitive positive interpretation J of $\mathfrak{D}$ in $\mathfrak{C}$. They are called* primitively positively bi-interpretable *if additionally $I \circ J$ and $J \circ I$ are pp-homotopic to the identity interpretation.*

Note that structures that are primitively positively bi-definable are in particular bi-interpretable (via 1-dimensional interpretations). It is known that two ω-categorical structures are primitively positively bi-interpretable if and only if there is a topological clone isomorphism between their polymorphism clones [42].

Example 5.2. The structure $(\mathbb{Q};<)$ and the structure $(\mathbb{I};\mathsf{m})$ from Example 2.7 are primitively positive bi-interpretable (Example 3.3.3 in [19]): the identity map I on $\mathbb{I}$ is a 2-dimensional primitive positive interpretation of $(\mathbb{I};\mathsf{m})$ in $(\mathbb{Q};<)$, and the projection $J\colon \mathbb{I} \to \mathbb{Q}$ to the first coordinate is a 1-dimensional interpretation of $(\mathbb{Q};<)$ in $(\mathbb{I};\mathsf{m})$.

Consider first the map I. The domain is $\mathbb{I} \subseteq \mathbb{Q}^2$, defined by a primitive positive formula $\delta(a,b) = a < b$. The relation

$$I^{-1}(\mathsf{m}) = \{((a,b),(c,d)) \mid I(a,b) \;\mathsf{m}\; I(c,d)\}$$

has a primitive positive definition over $(\mathbb{Q};<)$ by a formula $b = c$. The relation

$$I^{-1}(\equiv) = \{((a,b),(c,d)) \mid I(a,b) \;\equiv\; I(c,d)\}$$

[1] We follow the terminology from [2].

has a primitive positive definition $(a = c) \wedge (b = d)$. It follows that I is a primitive positive interpretation of $(\mathbb{I}; \mathsf{m})$ in $(\mathbb{Q}; <)$.

Concerning J, the domain is the full set $\mathbb{I}$ and hence defined by a formula $x = x$. The relation

$$J^{-1}(<) = \{(x, y) \mid J(x) < J(y)\} = \{((x_1, x_2), (y_1, y_2)) \mid x_1 < y_1\}$$

has a primitive positive definition

$$\exists u, v (u \mathsf{~m~} x) \wedge (u \mathsf{~m~} v) \wedge (v \mathsf{~m~} y).$$

The relation $J^{-1}(=)$ has a primitive positive definition $\phi(x, y)$ given by

$$\exists u (u \mathsf{~m~} x) \wedge (u \mathsf{~m~} y).$$

It follows that J is a primitive positive interpretation.

Observe that $J \circ I$ is pp-homotopic to the identity interpretation over $(\mathbb{Q}; <)$ since the relation

$$\{(a, b, c) \mid J \circ I(a, b) = c\}$$

has a primitive positive definition

$$a = c.$$

Finally, we show $I \circ J$ is pp-homotopic to the identity interpretation over $(\mathbb{I}; \mathsf{m})$. First note that the formula $\phi(x, y)$ is a primitive positive definition of the relation

$$\{(x, y) \mid x = (x_1, x_2), y = (y_1, y_2), x_1 = y_1\}$$

over $(\mathbb{I}; \mathsf{m})$. Observe that for $(x_1, x_2), (y_1, y_2)$ in the domain of $I \circ J$, we have $I \circ J((x_1, x_2), (y_1, y_2)) = I(J(x_1, x_2), J(y_1, y_2)) = (z_1, z_2)$ if and only if $x_1 = z_1$ and $y_1 = z_2$. Then

$$I \circ J(x, y) = z$$

is equivalent to the primitive positive formula

$$\phi(x, z) \wedge z \mathsf{~m~} y,$$

and hence $I \circ J$ is pp-homotopic to the identity interpretation. This finishes the proof that $(\mathbb{I}; \mathsf{m})$ and $(\mathbb{Q}; <)$ are bi-interpretable. We note that the proof of Theorem 6.7 generalizes this construction.

The proof of Lemma 3.4.1 in [19] shows in fact the following statement.

THEOREM 5.3. *Suppose $\mathfrak{D}$ has a primitive positive interpretation I in $\mathfrak{C}$, and $\mathfrak{C}$ has a primitive positive interpretation J in $\mathfrak{D}$ such that $J \circ I$ is pp-homotopic to the identity interpretation of $\mathfrak{C}$. Then for every first-order expansion $\mathfrak{C}'$ of $\mathfrak{C}$ there is a first-order expansion $\mathfrak{D}'$ of $\mathfrak{D}$ such that I is a primitive positive interpretation of $\mathfrak{D}'$ in $\mathfrak{C}'$ and J is a primitive positive interpretation of $\mathfrak{C}'$ in $\mathfrak{D}'$. The theorem is described in Figure 5.1.*

In particular, if $\mathfrak{C}$, $\mathfrak{D}$, $\mathfrak{C}'$ and $\mathfrak{D}'$ are as in Theorem 5.3, and $\mathfrak{C}'$ or $\mathfrak{D}'$ has a finite relational signature (and hence we may assume that both have a finite signature), then CSP($\mathfrak{C}'$) and CSP($\mathfrak{D}'$) have the same computational complexity (up to polynomial-time reductions) by Proposition 2.5. Now, let $\mathcal{C}$ and $\mathcal{D}$ denote the sets of first-order expansions of $\mathfrak{C}$ and $\mathfrak{D}$, respectively, and assume that the complexity of CSP($\mathfrak{D}$) is known for every $\mathfrak{D} \in \mathcal{D}$. It follows that we can deduce the complexity of CSP($\mathfrak{C}$) for every $\mathfrak{C} \in \mathcal{C}$.

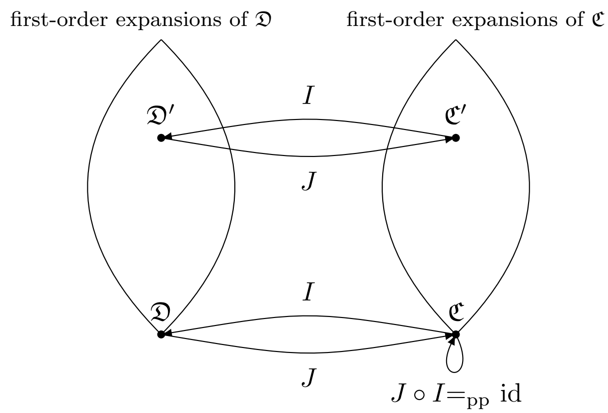

FIG. 5.1. *Visualisation of Theorem 5.3. We use the symbol $=_{\mathrm{pp}}$ to denote that two interpretations are pp-homotopic.*

TABLE 6.1
Overview of results and methods used for studying first-order expansions of the basic relations (unless otherwise stated). By 'Interval Algebra above s, f*' we mean a first-order expansion of* $(\mathbb{I}; \mathsf{s}, \mathsf{f})$.

Formalism	$(\mathbb{Q}^n; <_1, =_1, \ldots, <_n, =_n)$ classification	Classification transfer
Cardinal Direction Calculus	for $n = 2$	not used
Generalized CDC	for general n	not used
Interval Algebra	not used	used
Interval Algebra above s, f	for $n = 2$	used
Rectangle Algebra	for $n = 2$	used
n-dimensional Block Algebra	for general n	used

6. Applications. This section demonstrates that our dichotomy result for first-order expansions of the structure $(\mathbb{Q}^n; <_1, =_1, \ldots, <_n, =_n)$ from Section 4.5 can be combined with the complexity classification transfer result from Section 5 to obtain surprisingly strong new classification results. We obtain full classifications for the complexity of the CSP for first-order expansions of several interesting structures; an overview can be found in Table 6.1. Our results have stronger formulations when specialised to binary languages; this yields simple new proofs of known results (Sections 6.1 and 6.2.1) and solves long-standing open problems from the field of temporal and spatial reasoning (Sections 6.1 and 6.3). For the basic structures in Table 6.1, our results show that the CSP for a first-order expansion $\mathfrak{B}$ is polynomial-time solvable if and only if each relation in $\mathfrak{B}$ can be defined via an Ord-Horn formula.

6.1. Cardinal Direction Calculus. The *Cardinal Direction Calculus* (CDC) [76] is a formalism where the basic objects are the points in the plane, i.e., the domain is $\mathbb{Q}^2$. The basic relations correspond to eight cardinal directions (North, East, South, West and four intermediate ones) together with the equality relation. The basic relations can be viewed as pairs (R^1, R^2) for all choices of $R^1, R^2 \in \{<, =, >\}$, where each relation applies to the corresponding coordinate. The connection between cardinal directions and pairs (R^1, R^2) is described in Table 6.2. Let $\mathfrak{C}$ denote the structure containing the basic relations of CDC. The classical formulation of CDC contains all binary relations that are unions of relations in $\mathfrak{C}$. In the sequel, we will additionally be interested in the richer set of relations of arbitrary arity that are first-order definable in $\mathfrak{C}$.

THEOREM 6.1. *Let $\mathfrak{B}$ be a first-order expansion of $\mathfrak{C}$. Then exactly one of the following two cases applies.*

- *Each of $\theta_1(\mathrm{Pol}(\mathfrak{B}))$ and $\theta_2(\mathrm{Pol}(\mathfrak{B}))$ contains mi_3, min_3, mx_3, or ll_3, or one of their duals. In this case, $\mathrm{Pol}(\mathfrak{B})$ has a pwnu polymorphism. If the signature of $\mathfrak{B}$ is finite, then $\mathrm{CSP}(\mathfrak{B})$ is in P.*
- *$\mathrm{Pol}(\mathfrak{B})$ has a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$ and $\mathfrak{B}$ has a finite-signature reduct whose CSP is NP-complete.*

Proof. Clearly, every relation in $\mathfrak{B}$ is first-order definable in $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. Moreover, the relations in $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ are primitively positively definable in $\mathfrak{C}$:

- $x <_1 y$ is defined by $\exists z\big(x \text{ SW } z \wedge z \text{ NW } y\big)$,
- $x =_1 y$ is defined by $\exists z\big(x \text{ S } z \wedge z \text{ N } y\big)$, and
- the relations $<_2$ and $=_2$ can be defined analogously.

Thus, the result follows immediately from Theorem 4.35 because every first-order expansion of $\mathfrak{C}$ can be viewed as a first-order expansion of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. ☐

A slightly weaker version of the following result has been proved by Ligozat [76] using a fundamentally different approach.

COROLLARY 6.2. *Let $\mathfrak{B}$ be a binary first-order expansion of $\mathfrak{C}$. Then exactly one of the following cases applies.*

- *Each relation in $\mathfrak{B}$, viewed as a relation of arity four over $\mathbb{Q}$, has an Ord-Horn definition. In this case, $\mathfrak{B}$ has a pwnu polymorphism and $\mathrm{CSP}(\mathfrak{B})$ is in P.*
- *$\mathrm{Pol}(\mathfrak{B})$ has a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$ and $\mathrm{CSP}(\mathfrak{B})$ is NP-complete.*

Proof. Immediate consequence of Theorem 4.42. ☐

Ligozat [77] and Balbiani and Condotta [8] discuss the natural generalisation of CDC to the domain $\mathbb{Q}^n$: let CDC_n denote this generalisation. Balbiani and Condotta [8, Section 7] claim that a particular set of relations (referred to as *strongly preconvex*) is a maximal tractable subclass in CDC_3 that contains all basic relations, and they state that they have not been able to generalise this result to higher dimensions. Theorem 6.1 and Corollary 6.2 can immediately be generalised to this setting by using our results for $(\mathbb{Q}; <)^{(n)}$. We conclude that the Ord-Horn class is the unique maximal tractable subclass of CDC_n $(n \geq 2)$ that contains all basic relations.

6.2. Allen's Interval Algebra. We have already introduced Allen's Interval Algebra in Example 2.7 to illustrate interpretations. The complexity of all binary reducts of Allen's Interval Algebra have been classified in [72]. However, little is known about the complexity of the CSP for first-order reducts of Allen's Interval Algebra. In this section we obtain classification results for first-order expansions of some reducts of Allen's Interval Algebra. Recall the relations m ('meets'),

TABLE 6.2
The basic relations of Cardinal Direction Calculus.

=	N	E	S	W	NE	SE	SW	NW
$(=,=)$	$(=,>)$	$(>,=)$	$(=,<)$	$(<,=)$	$(>,>)$	$(>,<)$	$(<,<)$	$(<,>)$

s ('starts'), and f ('finishes') from Table 2.1. For reducts of Allen's Interval Algebra that contain m, these classification results are an immediate consequence of the transfer result from Section 5 (Section 6.2.1). For the first-order expansions of the structure that just contains the relations s and f (Section 6.2.2), we combine classification transfer with our classification for the first-order expansions of $(\mathbb{Q}; <_1, =_1, <_2, =_2)$ from Section 4.4.2.

6.2.1. First-order Expansions of $\{\mathsf{m}\}$. The following is a more explicit version of Theorem 3.4.3 in [19] (which only states that the CSP of a first-order expansions of $(\mathbb{I}; \mathsf{m})$ is polynomial-time solvable or NP-complete).

THEOREM 6.3. *Let $\mathfrak{B}$ be a first-order expansion of $(\mathbb{I}; \mathsf{m})$. Then exactly one of the following cases applies.*

- *The identity map on $\mathbb{I}$ is a 2-dimensional primitive positive interpretation of $\mathfrak{B}$ in $\mathfrak{U}$, $\mathfrak{X}$, $\mathfrak{I}$, or $\mathfrak{L}$. In this case, $\mathfrak{B}$ has a pwnu polymorphism, and if $\mathfrak{B}$ has a finite signature then* $\mathrm{CSP}(\mathfrak{B})$ *is polynomial-time solvable.*
- $\mathrm{Pol}(\mathfrak{B})$ *has a uniformly continuous minor-preserving map to* $\mathrm{Pol}(K_3)$ *(and $\mathfrak{B}$ has a finite-signature reduct whose CSP is NP-complete).*

Proof. From Example 5.2 we know that the structures $(\mathbb{I}; \mathsf{m})$ and $(\mathbb{Q}; <)$ are primitively positively bi-interpretable via interpretations I and J of dimension 2 and 1, respectively, where I is the identity map on $\mathbb{I}$. Theorem 5.3 implies that there exists a first-order expansion $\mathfrak{C}$ of $(\mathbb{Q}; <)$ such that I is a primitive positive interpretation of $\mathfrak{B}$ in $\mathfrak{C}$ and J is a primitive positive interpretation of $\mathfrak{C}$ in $\mathfrak{B}$.

If $\mathfrak{C}$ has a primitive positive interpretation in $\mathfrak{U}$, $\mathfrak{X}$, $\mathfrak{I}$, $\mathfrak{L}$, then $\mathfrak{B}$ has a primitive positive interpretation in one of those structures. Since each of $\mathfrak{U}$, $\mathfrak{X}$, $\mathfrak{I}$, and $\mathfrak{L}$ has a pwnu polymorphism (Theorem 4.2) and since the existence of pwnu polymorphisms is preserved by primitive positive interpretations (as was discussed in Section 2.3), the structure $\mathfrak{C}$ has a pwnu polymorphism. Furthermore, if $\mathfrak{B}$ has a finite signature, then $\mathrm{CSP}(\mathfrak{B})$ has a polynomial-time reduction to the CSP of one of those structures by Proposition 2.5. The polynomial-time tractability then follows from Theorem 4.2 and Theorem 4.3. Otherwise, Theorem 4.3 implies that $\mathrm{Pol}(\mathfrak{C})$ has a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$. Since there is also a uniformly continuous minor-preserving map from $\mathrm{Pol}(\mathfrak{B})$ to $\mathrm{Pol}(\mathfrak{C})$ by Theorem 2.11, we can compose maps and obtain a uniformly continous minor-preserving map from $\mathrm{Pol}(\mathfrak{B})$ to $\mathrm{Pol}(K_3)$. By Corollary 2.12, $\mathfrak{B}$ has a finite-signature reduct whose CSP is NP-complete. By Corollary 3.10, the two cases in the statement of the theorem are mutually exclusive. ∎

Nebel & Bürckert [87] proved (by a computer-generated proof) that if a reduct $\mathfrak{B}$ of Allen's Interval Algebra only contains relations that have an Ord-Horn definition when considered as a relation of arity four over $\mathbb{Q}$, then $\mathrm{CSP}(\mathfrak{B})$ is in P. Otherwise, and if it contains the relation m, it has an NP-hard CSP. Later on, Ligozat [75] presented a mathematical proof of this result. We can derive a stronger variant of the results by Nebel & Bürckert and Ligozat as a consequence of Theorem 6.3.

THEOREM 6.4. *Let $\mathfrak{B}$ be a binary first-order expansion of $(\mathbb{I}; \mathsf{m})$. Then exactly one of the following cases applies.*

- *Every relation of $\mathfrak{B}$, viewed as a relation of arity four over $\mathbb{Q}$, has an Ord-Horn definition. In this case, $\mathfrak{B}$ has a pwnu polymorphism and* $\mathrm{CSP}(\mathfrak{B})$ *is in P.*
- $\mathrm{Pol}(\mathfrak{B})$ *has a uniformly continuous minor-preserving map to* $\mathrm{Pol}(K_3)$ *and* $\mathrm{CSP}(\mathfrak{B})$ *is NP-complete.*

Proof. If $\mathrm{Pol}(\mathfrak{B})$ has a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$, then the statement follows from Corollary 2.12. Otherwise, Theorem 6.3 implies that $\mathfrak{B}$ has a pwnu polymorphism and the identity map I on $\mathbb{I}$ is a 2-dimensional primitive positive interpretation in $\mathfrak{U}$, $\mathfrak{X}$, $\mathfrak{I}$, or $\mathfrak{L}$. We claim that every relation R of $\mathfrak{B}$, considered as a relation of arity four over $\mathbb{Q}$, has an Ord-Horn definition. Let $\phi(u_1, u_2, v_1, v_2)$ be the first-order formula that defines $I^{-1}(R)$ over $(\mathbb{Q}; <)$. Note that if (u_1, u_2) is in the domain of I, then $u_1 < u_2$, so $\phi(u_1, u_2, v_1, v_2)$ implies $u_1 < u_2 \wedge v_1 < v_2$.

We first consider the case that $\mathfrak{B}$ has a 2-dimensional primitive positive interpretation in $\mathfrak{U}$, $\mathfrak{X}$, or $\mathfrak{I}$. In this case, ϕ is preserved by a pp-operation (Proposition 4.6), and we may assume that ϕ has the syntactic form described in Theorem 4.7. Since ϕ implies $u_1 < u_2 \wedge v_1 < v_2$, we may add these two conjuncts to ϕ; note that the resulting formula is still of the required form. We may additionally assume that ϕ is reduced, because every formula obtained from ϕ by removing literals is again of the required form. Each clause in ϕ has the form

$$y_1 \neq x \vee \cdots \vee y_k \neq x \vee z_1 \leq x \vee \cdots \vee z_l \leq x,$$

where all the variables are from $\{u_1, u_2, v_1, v_2\}$. We want to prove that ϕ is equivalent to an Ord-Horn formula, which we obtain by showing that $l \leq 1$ in every such clause.

If $u_1 \in \{z_1, \ldots, z_l\}$ and u_2 equals x then the conjunct $u_1 < u_2$ implies that the literal $u_1 \leq u_2$ is true in every satisfying assignment to ϕ, which means that the clause has no other literals by the assumption that ϕ is reduced, and we are done. If u_1 equals x and u_2 equals z_i, for some $i \in \{1, \ldots, l\}$, then the conjunct $u_1 < u_2$ implies that removing the literal $z_i \leq x$ would result in an equivalent formula, in contradiction to the assumption that ϕ is reduced. If u_1 equals z_i and u_2 equals z_j for some $i, j \in \{1, \ldots, l\}$, then the clause $u_1 < u_2$ implies that the literal $z_j \leq x$ is redundant, again in contradiction to the assumption that ϕ is reduced. We see that if one of u_1, u_2 is from $\{z_1, \ldots, z_l\}$, then the other variable cannot be from $\{x, z_1, \ldots, z_l\}$. The same argument applies to v_1 and v_2 and we conclude that $l \leq 1$.

Finally we consider the case that $\mathfrak{B}$ has a 2-dimensional primitive positive interpretation in $\mathfrak{L}$. In this case, Theorem 4.12 implies that every relation of $\mathfrak{B}$, considered as a relation of arity four over $\mathbb{Q}$, has a definition $\phi(u_1, u_2, v_1, v_2)$ by a conjunction of ll-Horn clauses

$$x_1 \neq y_1 \vee \cdots \vee x_m \neq y_m \vee z_1 < z_0 \vee \cdots \vee z_l < z_0 \vee (z_0 = z_1 = \cdots = z_l),$$

where the final disjunct might be missing and the variables x_i, y_i, z_i are from the set $\{u_1, u_2, v_1, v_2\}$. Again we may assume that ϕ contains the two clauses $u_1 < u_2$ and $v_1 < v_2$, and we may also assume that ϕ is reduced in the sense that whenever we remove a literal $z_i < z_0$ and remove z_i from the final disjunct $z_0 = z_1 = \cdots = z_l$, or if we remove the final disjunct entirely, we obtain a formula which is not equivalent to ϕ. It suffices to show that this implies that $l \leq 1$. Again, we break into cases. If both u_1 and u_2 are in $\{z_0, z_1, \ldots, z_l\}$, then the final disjunct is never satisfied, so we may assume that it is not present. If $u_1 \in \{z_1, \ldots, z_l\}$ and u_2 equals z_0, then the literal $z_i < z_0$ would be true in every satisfying assignment to ϕ, which means that the clause has no other literals by the assumption that ϕ is reduced, and we are done. If $u_1, u_2 \in \{z_1 \ldots, z_l\}$ we also obtain a contradiction to the assumption that ϕ is reduced, since $u_2 < z_0$ implies $u_1 < z_0$. If u_1 equals z_0 and u_2 equals z_i for some $i \in \{1, \ldots, l\}$, then the literal $z_i < z_0$ would be false and can be removed, in contradiction to ϕ being reduced. Again it follows that at most one of u_1 and u_2 can appear in $\{z_0, \ldots, z_l\}$. Analogous reasoning for v_1 and v_2 implies that $l \leq 1$.

The disjointness of the two cases follows from the disjointness of the cases in Theorem 6.3. ☐

6.2.2. First-order Expansions of $\{\mathsf{s}, \mathsf{f}\}$. Despite the obvious difference between the domains $\mathbb{I}$ and $\mathbb{Q}^2$, there is a way to use our classification of the first-order expansions of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$

to obtain classification results for first-order expansions of $(\mathbb{I}; \mathsf{s}, \mathsf{f})$. Our starting point is the following definability result.

LEMMA 6.5. $(\mathbb{I}; \mathsf{s}, \mathsf{f})$ *and* $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ *are primitively positively bi-definable.*

Proof. Consider the structure $(\mathbb{I}; <_1, =_1, <_2, =_2)$, where the relations $<_i$ and $=_i$ denote restrictions of $<_i$ and $=_i$ on $\mathbb{I} \subseteq \mathbb{Q}^2$. Note that the relation s can be defined over $(\mathbb{I}; <_1, =_1, <_2, =_2)$ by the formula $a =_1 b \wedge a <_2 b$. Similarly, the relation f can be defined over $(\mathbb{I}; <_1, =_1, <_2, =_2)$ by the formula $a =_2 b \wedge b <_1 a$. The relations in $(\mathbb{I}; <_1, =_1, <_2, =_2)$ can be primitively positively defined in $(\mathbb{I}; \mathsf{s}, \mathsf{f})$ as follows.

$$
\begin{aligned}
=_1 &= \big\{(a,b) \in \mathbb{I}^2 \mid \exists c\big(c \;\mathsf{s}\; a \wedge c \;\mathsf{s}\; b\big)\big\},\\
=_2 &= \big\{(a,b) \in \mathbb{I}^2 \mid \exists c\big(c \;\mathsf{f}\; a \wedge c \;\mathsf{f}\; b\big)\big\},\\
<_1 &= \big\{(a,b) \in \mathbb{I}^2 \mid \exists c,d\big(a =_1 c \wedge d \;\mathsf{f}\; c \wedge d \;\mathsf{s}\; b\big)\big\},\\
<_2 &= \big\{(a,b) \in \mathbb{I}^2 \mid \exists c,d\big(b =_2 c \wedge d \;\mathsf{s}\; c \wedge d \;\mathsf{f}\; a\big)\big\}.
\end{aligned}
$$

It follows that the structures $(\mathbb{I}; \mathsf{s}, \mathsf{f})$ and $(\mathbb{I}; <_1, =_1, <_2, =_2)$ are primitively positively interdefinable.

Claim. The structures $(\mathbb{I}; <_1, =_1, <_2, =_2)$ and $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ are isomorphic.

We prove the statement by a back-and-forth argument. Suppose that i is an isomorphism between a finite substructure $\mathfrak{A}$ of $(\mathbb{I}; <_1, =_1, <_2, =_2)$ and a finite substructure $\mathfrak{A}'$ of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. The sets A and A' denote (as usual) the domains of $\mathfrak{A}$ and $\mathfrak{A}'$, respectively. Let

$$
\begin{aligned}
A_1 &:= \{p \in \mathbb{Q} \mid \exists q\ (p,q) \in A\} \qquad & A_1' &:= \{p \in \mathbb{Q} \mid \exists q\ (p,q) \in A'\}\\
A_2 &:= \{q \in \mathbb{Q} \mid \exists p\ (p,q) \in A\} \qquad & A_2' &:= \{q \in \mathbb{Q} \mid \exists p\ (p,q) \in A'\}.
\end{aligned}
$$

Define $i_1 \colon A_1 \to A_1'$ by setting $i_1(p) = p'$ if there exist $q, q' \in \mathbb{Q}$ such that $i(p,q) = (p',q')$. Similarly, define $i_2 \colon A_2 \to A_2'$ by setting $i_2(q) = q'$ if there exist $p, p' \in \mathbb{Q}$ such that $i(p,q) = (p',q')$. Since i and i^{-1} are isomorphisms, i_1 and i_2 and their inverses are well-defined bijections and preserve $<$. By the homogeneity of $(\mathbb{Q}; <)$, there exist automorphisms α_1 and α_2 that extend i_1 and i_2.

For going forth, let $(a,b) \in \mathbb{I} \setminus A$. Then i is extended by setting $i(a,b) := (\alpha_1(a), \alpha_2(b))$. Since α_1 and α_2 preserve $<$, the extended map i preserves the relations $<_1, =_1, <_2, =_2$. The operations α_1 and α_2 are automorphisms of $(\mathbb{Q}; <)$, hence i is injective and i^{-1} preserves $<_1, =_1, <_2, =_2$. Therefore, the extension of i is an isomorphism.

For going back, let $(a', b') \in \mathbb{Q}^2 \setminus A'$. Then i is extended by setting $i(\alpha_1^{-1}(a'), \alpha_2^{-1}(b')) := (a', b')$; to prove that the extension is an isomorphism, we may argue similarly as in the forth step. Alternating between going back and going forth, we may thus construct an isomorphism between the two countable structures $(\mathbb{I}; <_1, =_1, <_2, =_2)$ and $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. ◇

The claim implies that $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ is isomorphic to the structure $(\mathbb{I}; <_1, =_1, <_2, =_2)$, which is primitively positive interdefinable with $(\mathbb{I}; \mathsf{s}, \mathsf{f})$, i.e., $(\mathbb{I}; \mathsf{s}, \mathsf{f})$ and $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ are primitively positively bi-definable. ∎

With the aid of Lemma 6.5, we can now proceed in a way that is similar to the proof of Theorem 6.3.

THEOREM 6.6. *Let* $\mathfrak{D}$ *be a first-order expansion of* $(\mathbb{I}; \mathsf{s}, \mathsf{f})$. *Then exactly one of the following cases applies.*

- $\mathfrak{D}$ *has a pwnu polymorphism. If* $\mathfrak{D}$ *has a finite relational signature, then* CSP($\mathfrak{D}$) *is in P.*

- $\mathrm{Pol}(\mathfrak{D})$ *has a uniformly continuous minor-preserving map to* $\mathrm{Pol}(K_3)$. *In this case,* $\mathfrak{D}$ *has a finite-signature reduct whose CSP is NP-complete.*

Proof. There is a first-order expansion $\mathfrak{C}$ of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ such that $\mathfrak{D}$ has a primitive positive interpretation in $\mathfrak{C}$ and vice versa, by Lemma 6.5 together with Theorem 5.3. If $\mathrm{Pol}(\mathfrak{D})$ has a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$, then $\mathfrak{D}$ has a finite-signature reduct whose CSP is NP-complete by Corollary 2.12. Otherwise, since there is a uniformly continuous minor-preserving map from $\mathrm{Pol}(\mathfrak{D})$ to $\mathrm{Pol}(\mathfrak{C})$ by Theorem 2.11, $\mathrm{Pol}(\mathfrak{C})$ does not have a uniformly continuous minor-preserving map to $\mathrm{Pol}(K_3)$ as well. By Theorem 4.35, $\mathfrak{C}$ has a pwnu polymorphism and $\mathrm{CSP}(\mathfrak{C})$ is in P if $\mathfrak{C}$ has a finite signature. Since primitive positive interpretations preserve the existence of pwnu polymorphisms, $\mathfrak{D}$ has a pwnu polymorphism as well. Moreover, if $\mathfrak{D}$ has a finite signature, then $\mathfrak{C}$ can be assumed to have a finite signature and $\mathrm{CSP}(\mathfrak{D})$ is in P by Proposition 2.5. By Corollary 3.10, the two cases in the statement of the theorem are mutually exclusive. □

We note that relations s and f are primitively positively definable in $\{\mathsf{m}\}$ but m is not primitively positively definable in $\{\mathsf{s}, \mathsf{f}\}$, so Theorem 6.6 is incomparable to Theorem 6.3.

6.3. Block Algebra. We will now study the *n-dimensional block algebra* ($\mathfrak{BA}_n$) by Balbiani et al. [10]. This formalism has become widespread since it can capture directional information in spatial reasoning, something that the classical RCC formalisms [91] cannot. Let $n \geq 1$ be an integer. The n-dimensional block algebra has the domain $\mathbb{I}^n$. For relations $R^1, \dots, R^n$ from the interval algebra, we write

$$\{((x_1, \dots, x_n), (y_1, \dots, y_n)) \in (\mathbb{I}^n)^2 \mid x_i \; R^i \; y_i, \; 1 \leq i \leq n\}.$$

The structure $\mathfrak{BA}_n$ contains all such relations, and we say that the relation $(R^1 | R^2 | \dots | R^n)$ is *basic* if $R^1, \dots, R^n$ are basic relations in the interval algebra. We note that $\mathfrak{BA}_1$ is the interval algebra and that $\mathfrak{BA}_2$ is often referred to as the *rectangle algebra* (RA) [63, 86].

We begin by studying first-order expansions of $(\mathbb{I}; \mathsf{m}) \boxtimes (\mathbb{I}; \mathsf{m}) = (\mathbb{I}^2; \mathsf{m}_1, =_1, \mathsf{m}_2, =_2)$. It is easy to see that the relation $=_i$ is primitively positively definable by m_i, hence it is equivalent to study the first-order expansions of the structure $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$. Note that the relation m_1 and m_2 over $\mathbb{I}^2$ can be written as $(\mathsf{m}|\top)$ and $(\top|\mathsf{m})$ respectively in the terminology of the Block Algebra. Also note that m_1 and m_2 are primitively positively definable over the basic relations of the rectangle algebra: for example, $\exists z (x \; (\mathsf{m}|\mathsf{p}) \; z \wedge y \; (\equiv|\mathsf{p}) \; z)$ is equivalent to $x \; (\mathsf{m}|\top) \; y$. The fact that every basic relation in the interval algebra has a primitive positive definition over $(\mathbb{I}; \mathsf{m})$ [4] now immediately implies that every RA relation has a primitive positive definition over $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$. Hence, the results below imply a classification of the Rectangle Algebra containing the basic relations.

Theorem 6.7. *Let* $\mathfrak{D}$ *be a first-order expansion of the structure* $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$. *Then there exists a first-order expansion* $\mathfrak{C}$ *of* $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ *such that* $\mathfrak{D}$ *has a 2-dimensional primitive positive interpretation in* $\mathfrak{C}$ *and* $\mathfrak{C}$ *has a 1-dimensional primitive positive interpretation in* $\mathfrak{D}$.

Furthermore, exactly one of the following two cases applies.

- $\mathfrak{D}$ *has a pwnu polymorphism. If the signature of* $\mathfrak{D}$ *is finite, then* $\mathrm{CSP}(\mathfrak{D})$ *is in P.*
- *There exists a uniformly continuous minor-preserving map from* $\mathrm{Pol}(\mathfrak{D})$ *to* $\mathrm{Pol}(K_3)$ *and* $\mathfrak{D}$ *has a finite-signature reduct whose CSP is NP-complete.*

Proof. For the first part of the statement we apply Theorem 5.3; so it suffices to prove that $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$ and $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ are primitively positively bi-interpretable via interpretations of dimension 2 and 1, respectively. This is basically the primitive positive bi-interpretation of $(\mathbb{I}; \mathsf{m})$ and $(\mathbb{Q}; <)$ from Example 3.3.3 in [19] performed in each dimension separately.

- There is a 2-dimensional interpretation I of $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$ in $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ whose domain $U \subseteq (\mathbb{Q}^2)^2$ has the primitive positive definition $\delta(a,b)$ given by $a <_1 b \wedge a <_2 b$. The interpretation $I \colon U \to \mathbb{I}^2$ is given by
$$((a_1, a_2), (b_1, b_2)) \mapsto ((a_1, b_1), (a_2, b_2)).$$
The relation
$$I^{-1}(\mathsf{m}_i) = \{((a,b),(c,d)) \mid I(a,b)\ \mathsf{m}_i\ I(c,d)\} \subseteq U^2$$
has the primitive positive definition $b =_i c$ in $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. The relation
$$I^{-1}((\equiv \mid \equiv)) = \{((a,b),(c,d)) \mid I(a,b)\ (\equiv \mid \equiv)\ I(c,d)\} \subseteq U^2$$
has the primitive positive definition $a =_1 c \wedge a =_2 c \wedge b =_1 d \wedge b =_2 d$ in $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$.
- There is a one-dimensional interpretation $J : \mathbb{I}^2 \to \mathbb{Q}^2$ of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ in $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$ given by
$$((p_1, p_2), (q_1, q_2)) \mapsto (p_1, q_1).$$
The relation
$$J^{-1}(<_i) = \{(x,y) \mid J(x) <_i J(y)\} \subseteq (\mathbb{I}^2)^2$$
has the primitive positive definition
$$\exists u, v (u\ \mathsf{m}_i\ x \wedge u\ \mathsf{m}_i\ v \wedge v\ \mathsf{m}_i\ y)$$
in $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$. The relation $J^{-1}(=_i)$ has the primitive positive definition
$$\phi_i(x,y) \coloneqq \exists u (u\ \mathsf{m}_i\ x \wedge u\ \mathsf{m}_i\ y)$$
in $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$ so $J^{-1}(=)$ has the definition $\phi_1(x,y) \wedge \phi_2(x,y)$.
- $J \circ I$ is pp-homotopic to the identity interpretation: We have
$$J(I(a,b)) = c \text{ if and only if } a = c.$$
- $I \circ J$ is pp-homotopic to the identity interpretation: Note that $\phi_i((x^1, x^2), (y^1, y^2))$ defines the relation
$$\{((x^1, x^2), (y^1, y^2)) \mid x^i =_1 y^i\} \subseteq (\mathbb{I}^2)^2$$
in $(\mathbb{I}; \mathsf{m}_i)$. If $x = ((x_1^1, x_2^1), (x_1^2, x_2^2))$ and $y = ((y_1^1, y_2^1), (y_1^2, y_2^2))$ are elements of $\mathbb{I}^2$, then
$$I(J(x), J(y)) = I((x_1^1, x_1^2), (y_1^1, y_1^2)) = ((x_1^1, y_1^1), (x_1^2, y_1^2)).$$
Therefore we have $I(J(x), J(y))\ (\equiv \mid \equiv)\ z$ if and only if
$$\phi_1(x,z) \wedge z\ \mathsf{m}_1\ y \wedge \phi_2(x,z) \wedge z\ \mathsf{m}_2\ y.$$

This concludes the proof of the first statement.

To prove the second statement, suppose that there is no uniformly continuous minor-preserving map from $\mathrm{Pol}(\mathfrak{D})$ to $\mathrm{Pol}(K_3)$—otherwise, we are done by Corollary 2.12. Since there is a primitive positive interpretation of $\mathfrak{C}$ in $\mathfrak{D}$, there is a uniformly continuous clone homomorphism from $\mathrm{Pol}(\mathfrak{D})$ to $\mathrm{Pol}(\mathfrak{C})$ by Lemma 2.10, and there cannot exist a uniformly continuous minor-preserving map from $\mathrm{Pol}(\mathfrak{C})$ to $\mathrm{Pol}(K_3)$. It now follows from Theorem 4.35 that $\mathrm{Pol}(\mathfrak{C})$ has a pwnu polymorphism and if $\mathfrak{C}$ has a finite relational signature then CSP($\mathfrak{C}$) is in P. By the first statement, there is also a primitive positive interpretation of $\mathfrak{D}$ in $\mathfrak{C}$. Therefore, the polynomial-time tractability of CSP($\mathfrak{D}$) in the finite signature case follows from Lemma 2.9. Moreover, there is a clone homomorphism from Pol($\mathfrak{C}$) to Pol($\mathfrak{D}$), again by Lemma 2.10. Therefore, $\mathfrak{D}$ has a pwnu polymorphism as well. By Corollary 3.10, the two cases in the statement of the theorem are mutually exclusive. ☐

We now consider binary first-order expansions of $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$. The proof combines arguments from Theorem 6.4 and Theorem 6.7.

THEOREM 6.8. *Let $\mathfrak{B}$ be a binary first-order expansion of $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$. Then exactly one of the following cases applies.*

- *Every relation of $\mathfrak{B}$, viewed as a relation of arity 8 over $\mathbb{Q}$, has an Ord-Horn definition. In this case, $\mathfrak{B}$ has a pwnu polymorphism and* CSP($\mathfrak{B}$) *is in P.*
- Pol($\mathfrak{B}$) *has a uniformly continuous minor-preserving map to* Pol(K_3) *and* CSP($\mathfrak{B}$) *is NP-complete.*

Proof. Let $\mathfrak{B}'$ be the first-order expansion of $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$ such that the structure $\mathfrak{B}$ has a 2-dimensional primitive positive interpretation I in $\mathfrak{B}'$ and $\mathfrak{B}'$ has a 1-dimensional primitive positive interpretation in $\mathfrak{B}$ which exists by Theorem 6.7. By Theorem 5.3, I may be taken to be the same interpretation $I\colon U \to \mathbb{I}^2$ as in the proof of Theorem 6.7, where $U = \{(a,b) \in \mathbb{Q}^2 \mid a <_1 b \wedge a <_2 b\}$ and

$$I : ((a_1, a_2), (b_1, b_2)) \mapsto ((a_1, b_1), (a_2, b_2)).$$

If Pol($\mathfrak{B}$) has a uniformly continuous minor-preserving map to Pol(K_3), then CSP($\mathfrak{B}$) is NP-hard by Corollary 2.12. Otherwise, it follows from Lemma 2.10 that Pol($\mathfrak{B}'$) does not have a uniformly continous minor-preserving map to Pol(K_3) and Theorem 4.35 implies that for each $i \in \{1,2\}$ there exists $f_i \in \text{Pol}(\mathfrak{B}')$ such that $\theta_i(f_i)$ equals min_3, mx_3, mi_3, ll_3, or one of their duals and that CSP($\mathfrak{B}'$) is in P. By Proposition 2.5, CSP($\mathfrak{B}$) is in P, too. In this case, $\mathfrak{B}$ has a pwnu polymorphism by Theorem 6.7; the theorem also implies that the two cases in the statement are mutually exclusive.

It remains to show that every (binary) relation of $\mathfrak{B}$, considered as a relation of arity 8 over $\mathbb{Q}$, has an Ord-Horn definition. Let R be a relation of $\mathfrak{B}$. Observe that it is sufficient to show that the 4-ary relation $I^{-1}(R)$ has a definition ϕ that is a conjunction of clauses of the form

$$x_1 \neq_{i_1} y_1 \vee \cdots \vee x_m \neq_{i_m} y_m \vee z_1 \circ z_0, \tag{6.1}$$

where $i_j \in \{1,2\}$, $\circ \in \{<_1, \leq_1, =_1, <_2, \leq_2, =_2\}$, it is permitted that $m = 0$ and the last disjunct may be omitted. With this in mind, ve(ϕ) (as defined just before Proposition 4.32) is the desired Ord-Horn definition of R viewed as a relation of arity 8 over $\mathbb{Q}$. By Lemma 4.20, we may focus on the situation when $\theta_i(f_i) \in \{\text{min}_3, \text{mx}_3, \text{mi}_3, \text{ll}_3\}$ (since an Ord-Horn definition with reversed ordering in one of the dimensions results in an Ord-Horn definition again).

Let $\phi(u_1, u_2, v_1, v_2)$ be the first-order definition of $I^{-1}(R)$ over $(\mathbb{Q}^2; <_1, =_1, <_2, =_2)$. By the definition of I, if $I(u_1, u_2) = u$, then $u_1 <_1 u_2 \wedge u_1 <_2 u_2$ and if $I(v_1, v_2) = v$, then $v_1 <_1 v_2 \wedge v_1 <_2 v_2$. Therefore, ϕ implies that the four conjuncts above hold.

Suppose first that $\theta_i(f_i) = \text{ll}_3$, $i = 1, 2$, then, by Proposition 3.5 and Theorem 4.12, $\theta_i(\text{Pol}(\mathfrak{B}'))$ contains an ll-operation for both i. By Proposition 4.32, we may assume that ϕ is a conjunction of clauses of the form

$$x_1 \neq_{i_1} y_1 \vee \cdots \vee x_m \neq_{i_m} y_m \vee z_1 <_j z_0 \vee \cdots \vee z_\ell <_j z_0 \vee (z_0 =_j z_1 =_j \cdots =_j z_\ell)$$

for $i_1, \dots, i_m, j \in \{1,2\}$ and where the last disjunct may be omitted. Since ϕ implies $u_1 <_1 u_2$, $u_1 <_2 u_2$, $v_1 <_1 v_2$ and $v_1 <_2 v_2$, we may add these conjuncts to ϕ without loss of generality; note that the formula is still of the required form. By an analogous argument as in the proof of Theorem 6.4, we can assume that ϕ is of the form (6.1).

Next, let $\theta_1(f_1) \in \{\mathrm{min}_3, \mathrm{mx}_3, \mathrm{mi}_3\}$ and $\theta_2(f_2) = \mathrm{ll}_3$. As in the previous paragraph, $\theta_2(\mathrm{Pol}(\mathfrak{B}'))$ contains an ll-operation. Similarly, by Proposition 3.5 and Proposition 4.6, $\theta_1(\mathrm{Pol}(\mathfrak{B}'))$ contains a pp-operation. By Proposition 4.30, R may be defined by a conjunction of weakly 1-determined clauses of the form

$$u_1 \neq_2 v_1 \vee \cdots u_m \neq_2 v_m \vee y_1 \neq_1 x \vee \cdots \vee y_k \neq_1 x \vee z_1 \leq_1 x \vee \cdots \vee z_l \leq_1 x$$

together with 2-determined clauses of the form

$$x_1 \neq_2 y_1 \vee \cdots \vee x_m \neq_2 y_m \vee z_1 <_2 z_0 \vee \cdots \vee z_\ell <_2 z_0 \vee (z_0 =_2 z_1 =_2 \cdots =_2 z_\ell).$$

Again, we may add the implied conjuncts $u_1 <_1 u_2$, $u_1 <_2 u_2$, $v_1 <_1 v_2$ and $v_1 <_2 v_2$ to the defining formula. Now we may use the same arguments as in the proof of Theorem 6.4 to prove that each of the clauses may be taken to be a clause of the form (6.1). The proof with the two dimensions exchanged is analogous.

Finally, assume that $\theta_i(f_i) \in \{\mathrm{min}_3, \mathrm{mx}_3, \mathrm{mi}_3\}$ for $i \in \{1,2\}$. We can reason analogously to the previous paragraph and conclude that $\theta_i(\mathrm{Pol}(\mathfrak{B}'))$, $i \in \{1,2\}$, contains a pp-operation. Furthermore, Theorem 4.12 implies that we may assume that $\theta_i(\mathrm{Pol}(\mathfrak{B}'))$, $i \in \{1,2\}$, does not contain a lex-operation (otherwise it would contain an ll-operation and this case would be covered by one of the previous cases). By applying Proposition 4.31 twice (the second time with the roles of the dimensions exchanged), every relation of $\mathfrak{B}'$ can be defined by a formula $\psi_1 \wedge \psi_2$ such that ψ_i is a conjunction of clauses of the form

$$y_1 \neq_i x \vee \cdots \vee y_k \neq_i x \vee z_1 \leq_i x \vee \cdots \vee z_l \leq_i x.$$

As in the previous cases, we may add the conjuncts $u_1 <_1 u_2$, $u_1 <_2 u_2$, $v_1 <_1 v_2$ and $v_1 <_2 v_2$ to the formula $\psi_1 \wedge \psi_2$ without loss of generality. Now we may proceed analogously to the proof of Theorem 6.4 and show that if the formula is reduced, it is of the form (6.1). ◻

Balbiani et al. [9] have presented a tractable subclass—consisting of the so-called *strongly preconvex* relations—of the rectangle algebra. They write the following on page 447.

> The subclass generated by the set of the strongly preconvex relations is now the biggest known tractable set of RA which contains the 169 atomic relations. An open question is: is this subclass a maximal tractable subclass which contains the atomic relations?

We answer their question affirmatively. We know that every basic RA relation has a primitive positive definition in $(\mathbb{I}^2; \mathsf{m}_1, \mathsf{m}_2)$, and we observed in the beginning of this section that the relation m_1 is primitively positively definable with the aid of the basic relations $(\mathsf{m}|\mathsf{p})$ and $(\equiv|\mathsf{p})$, and m_2 is analogously primitively positively definable from $(\mathsf{p}|\mathsf{m})$ and $(\mathsf{p}|\equiv)$. We have thus proved (via Theorem 6.8) that the Rectangle Algebra contains a single maximal subclass that is polynomial-time solvable and contains all basic relations. The relations in this subclass are definable via Ord-Horn formulas, and Balbiani et al. [10, Section 6.2] have proved that strongly preconvex relations coincide with Ord-Horn-definable relations.

We continue by analysing the n-dimensional block algebra when $n > 2$. The approach is similar to the approach used for the rectangle algebra: we use complexity transfer to deduce a classification for first-order expansions of $(\mathbb{I}^n; \mathsf{m}_1, \dots, \mathsf{m}_n)$ from the classification for first-order expansions of $(\mathbb{Q}^n; <_1, =_1, \dots, <_n, =_n)$. The relations $\mathsf{m}_1, \dots, \mathsf{m}_n$ are primitively positively definable from the basic relations of the n-dimensional block algebra, so we obtain in particular a classification of the complexity of the CSP for all first-order expansions of the basic relations of the n-dimensional block algebra.

THEOREM 6.9. *Let $\mathfrak{D}$ be a first-order expansion of the structure $(\mathbb{I}^n; \mathsf{m}_1, \dots, \mathsf{m}_n)$. Then there exists a first-order expansion $\mathfrak{C}$ of $(\mathbb{Q}^n; <_1, =_1, \dots, <_n, =_n)$ such that $\mathfrak{D}$ has a 2-dimensional primitive positive interpretation in $\mathfrak{C}$ and $\mathfrak{C}$ has a 1-dimensional primitive positive interpretation in $\mathfrak{D}$. Furthermore, exactly one of the following two cases applies.*

- *$\mathfrak{D}$ has a pwnu polymorphism. If the signature of $\mathfrak{D}$ is finite, then* CSP($\mathfrak{D}$) *is in P.*
- *There exists a uniformly continuous minor-preserving map from* Pol($\mathfrak{D}$) *to* Pol(K_3) *and $\mathfrak{D}$ has a finite-signature reduct whose* CSP *is NP-complete.*

Proof. Straightforward generalisation of Theorem 6.7. □

It has been known for a long time that the set of Ord-Horn-definable relations is a tractable fragment of the n-dimensional Block Algebra [10]. In that article (pp. 907–908), Balbiani et al. note the following.

> The problem of the maximality of this tractable subset [Ord-Horn] remains an open problem. Usually to prove the maximality of a fragment of a relational algebra an extensive machine-generated analysis is used. Because of the huge size ... we cannot proceed in the same way.

We answer this question in the affirmative: the subset of relations in $\mathfrak{B}\mathfrak{A}_n$ that can be viewed as arity-$4n$ relations with an Ord-Horn definition, is a maximal tractable subclass. Furthermore, it is the only maximal subclass that is tractable and contains all basic relations. To see this, we proceed in the same way as in the analysis of the Rectangle Algebra. First of all, every basic relation in $\mathfrak{B}\mathfrak{A}_n$ has a primitive positive definition in $(\mathbb{I}^n; \mathsf{m}_1, \dots, \mathsf{m}_n)$, and the relations $\mathsf{m}_1, \dots, \mathsf{m}_n$ are easily seen to be primitively positive definable in the basic relations of $\mathfrak{B}\mathfrak{A}_n$. It follows from the corollary below that the only maximal subclass of $\mathfrak{B}\mathfrak{A}_n$ that is polynomial-time solvable and contains all basic relations is the Ord-Horn class.

COROLLARY 6.10. *Let $\mathfrak{B}$ be a binary first-order expansion of $(\mathbb{I}^n; \mathsf{m}_1, \dots, \mathsf{m}_n)$. Then exactly one of the following cases applies.*

- *Each relation in $\mathfrak{B}$, viewed as a relation of arity $4n$ over $\mathbb{Q}$, has an Ord-Horn definition. In this case, $\mathfrak{B}$ has a pwnu polymorphism and* CSP($\mathfrak{B}$) *is in P.*
- Pol($\mathfrak{B}$) *has a uniformly continuous minor-preserving map to* Pol(K_3) *and $\mathfrak{B}$ has a finite-signature reduct whose* CSP *is NP-complete.*

Proof. Recall Theorem 6.9 and let $\mathfrak{B}'$ be the first-order expansion of $(\mathbb{Q}^n; <_1, =_1, \dots, <_n, =_n)$ such that $\mathfrak{B}$ has a 2-dimensional primitive positive interpretation I in $\mathfrak{B}'$ and $\mathfrak{B}'$ has a 1-dimensional primitive positive interpretation in $\mathfrak{B}$. By combining Corollary 2.12, Lemma 2.10 and Theorem 4.40, we deduce by the same argumentation as in the proof of Theorem 6.8 that if the statement in the second item does not hold, then $\mathfrak{B}'$ and $\mathfrak{B}$ have a pwnu polymorphism and CSP($\mathfrak{B}'$) and CSP($\mathfrak{B}$) are in P. Moreover, in this case, Theorem 4.40 implies that for every $p \in \{1, \dots, n\}$, there is an $f_p \in \mathrm{Pol}(\mathfrak{B}')$ such that $\theta_p(f_p)$ equals min_3, mx_3, mi_3, ll_3, or one of their duals.

It remains to show that if CSP($\mathfrak{B}$) is in P, then every (binary) relation of $\mathfrak{B}$, considered as a relation of arity $4n$ over $\mathbb{Q}$, has an Ord-Horn definition. Let R be a relation of $\mathfrak{B}$. Observe that, as in the proof of Theorem 6.8, it is enough to show that the 4-ary relation $I^{-1}(R)$ has a definition

$\phi(u_1, u_2, v_1, v_2)$ that is a conjunction of clauses of the form

$$x_1 \neq_{i_1} y_1 \vee \cdots \vee x_m \neq_{i_m} y_m \vee z_1 \circ z_0, \tag{6.2}$$

where $i_j \in \{1, \ldots, n\}$, $\circ \in \{<_1, \leq_1, =_1, \ldots, <_n, \leq_n, =_n\}$, it is permitted that $m = 0$ and the last disjunct may be omitted; then $\text{ve}(\phi)$ will be the desired Ord-Horn definition of R viewed as a relation of arity $4n$ over $\mathbb{Q}$.

By Lemma 4.20, we may focus on the situation that $f_p \in \{\text{min}_3, \text{mx}_3, \text{mi}_3, \text{ll}_3\}$ for every $p \in \{1, \ldots, n\}$. Therefore, by Proposition 4.6 and Theorem 4.12, there is a set $S \subseteq \{1, \ldots, n\}$ such that, for each $p \in S$, $\theta_p(\text{Pol}(\mathfrak{B}'))$ contains an ll-operation and, for each $p \in \{1, \ldots, n\} \setminus S$, $\theta_p(\text{Pol}(\mathfrak{B}'))$ contains a pp-operation but not a lex-operation. We may therefore assume that ϕ has the syntactic form described in Proposition 4.38. Moreover, we may assume that ϕ contains the conjuncts $u_1 <_j u_2$ and $v_1 <_j v_2$, $j = 1, \ldots, n$, since these are implied by ϕ (see the proof of Theorem 6.8 for more details). For each of the two types of clauses that appear in ϕ, we may use the same case distinction as in Theorems 6.4 and 6.8 to show that each of the clauses is of the form (6.2). This concludes the proof. ◻

Balbiani et al. [10, p. 908] also raise the following question:

> ...the question also arises as to how the qualitative constraints [Block Algebra] we have been considering could be integrated into a more general setting to include metric constraints.

If one focuses on tractable subclasses that contain all basic relations, then such an integration is indeed possible. Since our results imply that the relations in such a tractable subclass must be definable via Ord-Horn formulas, they can immediately be embedded into the metric framework suggested by Jonsson and Bäckström [67] and Koubarakis [71]. Under the same assumptions, this holds for the cardinal direction calculus and Allen's Interval Algebra, too.

7. Conclusions and Open Problems.

We proved that the CSPs for first-order expansions of $(\mathbb{Q}; <)^{(n)}$ satisfy a complexity dichotomy: they are in P or NP-complete. Using a general complexity transfer method, we prove that first-order expansions of the basic relations of the cardinal direction calculus, Allen's Interval Algebra, and the n-dimensional block algebra have a CSP complexity dichotomy. Less obviously, the complexity transfer method can also be applied to show that first-order expansions of the relations s and f of Allen's Interval Algebra have a complexity dichotomy. All of the results can be specialised for binary signatures, in which case we obtain new and conceptually simple proofs of results that have first been shown with the help of a computer (Theorem 6.4) or that answer several questions from the literature (Section 6.3).

Our results also imply that the so-called *meta-problem* of complexity classification is decidable: given finitely many first-order formulas that define a first-order expansion $\mathfrak{D}$ of one of the structures for which we obtained a complexity classification, one can effectively decide whether $\text{CSP}(\mathfrak{D})$ is in P or NP-complete. This follows from the general fact that for homogeneous finitely bounded structures that are model-complete cores and have an extremely amenable automorphism group (all of these assumptions are satisfied by our structures) the condition of the tractability conjecture (see Corollary 2.12) can be decided effectively (essentially by checking exhaustively for the existence of a *diagonally canonical pseudo Siggers polymorphism*); since these results are not new we refer to [19, Section 11.6] for details.

One may wonder about first-order reducts of $(\mathbb{Q}; <)^{(n)}$ rather than just first-order expansions. Classifying the complexity of the CSP for this class of structures will be a challenging project. It is easy to see that every CSP for a finite-domain structure can be formulated in this way. Also all first-order reducts of the infinite Johnson graphs $J(\omega, n)$ fall into this class (for example, the line graph of the countably infinite clique when $n = 2$). Such a project would also include the following interesting classification problem.

The *age* of a relational structure $\mathfrak{B}$ is the class of all finite structures that embed into $\mathfrak{B}$. It follows from Fraïssé's theorem (or by a direct back-and-forth argument) that two homogeneous structures with the same age are isomorphic (see, e.g. [65, Theorem 6.1.2]). Let $\mathfrak{A}_1$ and $\mathfrak{A}_2$ be homogeneous structures with disjoint relational signatures τ_1 and τ_2 and without algebraicity (see [65]; the structure $(\mathbb{Q}; <)$ is an example of such a structure without algebraicity). It is well known that there exists an up to isomorphism unique countable homogeneous $(\tau_1 \cup \tau_2)$-structure whose age consists of all structures whose τ_1-reduct is in the age of $\mathfrak{A}_1$ and whose τ_2-reduct is in the age of $\mathfrak{A}_2$; this structure is called the *generic combination* of $\mathfrak{A}_1$ and $\mathfrak{A}_2$, and will be denoted by $\mathfrak{A}_1 * \mathfrak{A}_2$. It can be shown by a back-and-forth argument that the τ_1-reduct of $\mathfrak{A}_1 * \mathfrak{A}_2$ is isomorphic to $\mathfrak{A}_1$ and the τ_2-reduct is isomorphic to $\mathfrak{A}_2$. The notion of generic combinations can be defined also for ω-categorical structures without algebraicity [23]: $\mathfrak{A}_1 * \mathfrak{A}_2$ is then defined as the $(\tau_1 \cup \tau_2)$-reduct of the generic combination of a homogeneous expansion of $\mathfrak{A}_1$ and a homogeneous expansion of $\mathfrak{A}_2$ (it can be shown that this is well-defined). If $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are first-order reducts of $(\mathbb{Q}; <)$ then the complexity of $\mathrm{CSP}(\mathfrak{A}_1 * \mathfrak{A}_2)$ has been classified recently [24]. A complexity classification of $\mathrm{CSP}(\mathfrak{B})$ for first-order reducts $\mathfrak{B}$ of $(\mathbb{Q}; <) * (\mathbb{Q}; <)$, however, is open.

Proposition 7.1. *For every first-order reduct $\mathfrak{B}$ of $(\mathbb{Q}; <) * (\mathbb{Q}; <)$, there exists a first-order reduct $\mathfrak{C}$ of $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$ such that $\mathfrak{C}$ is homomorphically equivalent to $\mathfrak{B}$.*

Proof. Let $\mathfrak{A}$ be the $\{<_1, <_2\}$-reduct of $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$. For $i \in \{1, 2\}$ there exists an isomorphism α_i between the $\{<_i\}$-reduct of $(\mathbb{Q}; <) * (\mathbb{Q}; <)$ and $(\mathbb{Q}; <)$. Then $e \colon d \mapsto (\alpha_1(d), \alpha_2(d))$ is an embedding of $(\mathbb{Q}; <) * (\mathbb{Q}; <)$ into $\mathfrak{A}$. Conversely, if we fix a linear extension of $<_1$ and $<_2$ in $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$, then the $\{<_1, <_2\}$-reduct of the resulting structure embeds into $(\mathbb{Q}; <) * (\mathbb{Q}; <)$ (see, e.g., [19, Lemma 4.1.7]). This shows that $\mathfrak{A}$ has an injective homomorphism h to $(\mathbb{Q}; <) * (\mathbb{Q}; <)$.

Let $\mathfrak{B}$ be a first-order reduct of $(\mathbb{Q}; <) * (\mathbb{Q}; <)$. Every first-order formula ϕ over $(\mathbb{Q}; <) * (\mathbb{Q}; <)$ is equivalent to a quantifier-free formula in conjunctive normal form, because $(\mathbb{Q}; <) * (\mathbb{Q}; <)$ is homogeneous and ω-categorical. Replace each atomic subformula of ϕ of the form $\neg(x <_i y)$, for $i \in \{1, 2\}$, by $y <_i x \vee x = y$. Then replace each subformula of the form $x \neq y$ by $x <_1 y \vee y <_1 x$. The resulting formula is equivalent over $(\mathbb{Q}; <) * (\mathbb{Q}; <)$. Each formula that defines a relation of $\mathfrak{B}$ and is written in this form can be interpreted over $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$ instead of $(\mathbb{Q}; <) * (\mathbb{Q}; <)$; let $\mathfrak{C}$ be the obtained first-order reduct of $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$.

Since the first-order definitions of the relations of $\mathfrak{B}$ are quantifier-free, the embedding e of $(\mathbb{Q}; <) * (\mathbb{Q}; <)$ into $\mathfrak{A}$ is also an embedding of $\mathfrak{B}$ into $\mathfrak{C}$. We claim that h is a homomorphism from $\mathfrak{C}$ to $\mathfrak{B}$. This follows from the fact that h is a homomorphism from $(\mathbb{Q}; <) \boxtimes (\mathbb{Q}; <)$ to $\mathfrak{A}$ and that the defining formulas for the relations of $\mathfrak{B}$ and $\mathfrak{C}$ do not involve negation. ◻

Another way forward is to study basic structures other than $(\mathbb{Q}; <)$. Here, temporal reasoning is a source of examples with applications in, for instance, AI. An important time model used in temporal reasoning is *branching time*, where for every point in time the past is linearly ordered, but the future is partially ordered. This motivates the so-called *left-linear point algebra* [54, 64], which is a relation algebra with four basic relations, denoted by $=$, $<$, $>$, and $|$. Here, $x|y$ signifies that x and y are incomparable in time, and '$x < y$' signifies that x is earlier in time than y. The branching-time satisfiability problem can be formulated as $\mathrm{CSP}(\mathfrak{B})$ for an ω-categorical structure $\mathfrak{B}$ [38]. One possible concrete description of the structure $\mathfrak{B}$, described by Adeleke and Neumann [1], is to let $\mathfrak{B} = (B; <, |, =)$ where B is the set of finite sequences of rational numbers. For arbitrary

$a = (a_1, \ldots, a_n)$ and $b = (b_1, \ldots, b_m)$ in B with $n \leq m$, $a < b$ holds if one of the following conditions hold:

1. $n < m$ and $a_i = b_i$ for $1 \leq i \leq n$, or
2. $a_i = b_i$ for $1 \leq i < n$ and $a_n < b_n$.

The remaining relations are defined in the obvious way.

Branching time has been used, for example, in automated planning [52], as the basis for temporal logics [57], and as the basis for a generalisation of Allen's Interval Algebra [90]. In particular, the complexity of the branching variant[2] of Allen's Interval Algebra has recently gained attraction [16, 17, 15, 60, 55]. Some complexity results for the branching interval algebra $\mathfrak{B}$ are presented in these publications, but the big picture is missing, even for first-order expansions of the basic relations. Our classification transfer result (Theorem 5.3) is applicable to this problem, but since there is currently no complexity classification of the CSPs for first-order reducts of $\mathfrak{B}$, we cannot present a full classification for CSPs of first-order expansions of the basic branching interval relations. Naturally, there is no polymorphism-based description of the tractable fragments either, so we cannot analyse the complexity of CSPs for first-order expansions of the structure $(B^n; <_1, |_1, =_1, \ldots, <_n, |_n, =_n)$ in the style of Theorem 4.40. However, given a complexity classification of the CSPs for first-order expansions of $\mathfrak{B}$ in place, then Theorem 5.3 is immediately applicable, and we do not see any fundamental problem that prevents us from generalising Theorem 4.40 to expansions of $(B^n; <_1, |_1, =_1, \ldots, <_n, |_n, =_n)$ as long as the tractable fragments can be described via polymorphisms and nice syntactic normal forms.

A related time model encountered in computer science is *partially ordered time* (po-time). This model has various applications in, for instance, the analysis of concurrent and distributed systems [6, 74]. In po-time, both the past and the future of a time point are partially ordered. This implies that time becomes a partial order with four basic relations $=$, $<$, $>$ and $|$, signifying "equal", "before", "after" and "unrelated", respectively. The satisfiability problem for po-time can convieniently be formulated with the *random partial order* $(P; <)$, and Kompatscher and Van Pham [70] have presented a full complexity classification of the CSP for all first-order reducts of the random partial order. Combined with Theorem 5.3, this gives us a full classification of the CSP for first-order expansions of the basic relations of the po-time analogue of the interval algebra. This generalisation of the interval algebra has been studied by Zapata et al. [97]. Kompatscher and Van Pham describe the tractable fragments of the random partial order with the aid of polymorphisms. Hence, it seems conceivable that their result can be generalised to a complexity classification of the CSP for first-order expansions of $(P^n; <_1, |_1, =_1, \ldots, <_n, |_n, =_n)$ by utilising the ideas behind the proof of Theorem 4.40.

Acknowledgements. We thank Johannes Greiner for his comments on an earlier version of this article and Jakub Rydval for letting us use his picture in Figure 4.1.

REFERENCES

[1] S. A. Adeleke and P. M. Neumann, *Relations related to betweenness: their structure and automorphisms*, Memoirs of the American Mathematical Society, 131 (1998), pp. viii+125.

[2] G. Ahlbrandt and M. Ziegler, *Quasi-finitely axiomatizable totally categorical theories*, Annals of Pure and Applied Logic, 30 (1986), pp. 63–82.

[2] Various ways of defining the formalism are possible [90]. We restrict our attention to the most well-known end-point-based formalism that contains 19 basic relations.

[3] J. F. Allen, *Maintaining knowledge about temporal intervals*, Communications of the ACM, 26 (1983), pp. 832–843.

[4] J. F. Allen and P. J. Hayes, *A common-sense theory of time*, in Proc. 9th International Joint Conference on Artificial Intelligence (IJCAI-1985), 1985, pp. 528–531.

[5] J. F. Allen and J. A. G. M. Koomen, *Planning using a temporal world model*, in Proc. 8th International Joint Conference on Artificial Intelligence (IJCAI-1983), 1983, pp. 741–747.

[6] F. D. Anger, *On Lamport's interprocessor communication model*, ACM Transactions on Programming Languages and Systems, 11 (1989), pp. 404–417.

[7] F. Baader and J. Rydval, *Description logics with concrete domains and general concept inclusions revisited*, in Proc. 10th International Joint Conference on Automated Reasoning (IJCAR-2020), 2020, pp. 413–431.

[8] P. Balbiani and J. Condotta, *Spatial reasoning about points in a multidimensional setting*, Applied Intelligence, 17 (2002), pp. 221–238.

[9] P. Balbiani, J. Condotta, and L. F. del Cerro, *A new tractable subclass of the rectangle algebra*, in Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-1999), 1999, pp. 442–447.

[10] P. Balbiani, J.-F. Condotta, and L. F. del Cerro, *Tractability results in the block algebra*, Journal of Logic and Computation, 12 (2002), pp. 885–909.

[11] L. Barto, B. Bodor, M. Kozik, A. Mottet, and M. Pinsker, *Symmetries of graphs and structures that fail to interpret a finite thing*, in LICS, 2023, pp. 1–13, https://doi.org/10.1109/LICS56636.2023.10175732.

[12] L. Barto, M. Kompatscher, M. Olšák, T. V. Pham, and M. Pinsker, *Equations in oligomorphic clones and the constraint satisfaction problem for ω-categorical structures*, Journal of Mathematical Logic, 19 (2019), p. #1950010.

[13] L. Barto, J. Opršal, and M. Pinsker, *The wonderland of reflections*, Israel Journal of Mathematics, 223 (2018), pp. 363–398.

[14] L. Barto and M. Pinsker, *The algebraic dichotomy conjecture for infinite domain constraint satisfaction problems*, in Proc. 31th Annual IEEE Symposium on Logic in Computer Science (LICS-2016), 2016, pp. 615–622.

[15] A. Bertagnon, M. Gavanelli, A. Passantino, G. Sciavicco, and S. Trevisani, *The Horn fragment of branching algebra*, in Proc. 27th International Symposium on Temporal Representation and Reasoning (TIME-2020), 2020, pp. 5:1–5:16.

[16] A. Bertagnon, M. Gavanelli, A. Passantino, G. Sciavicco, and S. Trevisani, *Branching interval algebra: An almost complete picture*, Information and Computation, 281 (2021), p. 104809.

[17] A. Bertagnon, M. Gavanelli, G. Sciavicco, and S. Trevisani, *On (maximal, tractable) fragments of the branching algebra*, in Proc. 35th Italian Conference on Computational Logic (CILC-2020), 2020, pp. 113–126.

[18] M. Bienvenu, B. ten Cate, C. Lutz, and F. Wolter, *Ontology-based data access: A study through disjunctive Datalog, CSP, and MMSNP*, ACM Transactions of Database Systems, 39 (2014), p. 33.

[19] M. Bodirsky, *Complexity of Infinite-Domain Constraint Satisfaction*, Cambridge University Press, 2021.

[20] M. Bodirsky and B. Bodor, *Canonical polymorphisms of Ramsey structures and the unique interpolation property*, in Proc. 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS-2021), 2021, pp. 1–13.

[21] M. Bodirsky, H. Chen, and M. Pinsker, *The reducts of equality up to primitive positive interdefinability*, Journal of Symbolic Logic, 75 (2010), pp. 1249–1292.

[22] M. Bodirsky, H. Chen, and M. Wrona, *Tractability of quantified temporal constraints to the max*, International Journal of Algebra and Computation, 24 (2014), pp. 1141–1156.

[23] M. Bodirsky and J. Greiner, *The complexity of combinations of qualitative constraint satisfaction problems*, Logical Methods in Computer Science, 16 (2020).

[24] M. Bodirsky, J. Greiner, and J. Rydval, *Tractable combinations of temporal CSPs*, CoRR, abs/2012.05682 (2020).

[25] M. Bodirsky and M. Grohe, *Non-dichotomies in constraint satisfaction complexity*, in Proc. 35th International Colloquium on Automata, Languages and Programming (ICALP-2008), 2008, pp. 184 –196.

[26] M. Bodirsky and P. Jonsson, *A model-theoretic view on qualitative constraint reasoning*, Journal of Artificial Intelligence Research, 58 (2017), pp. 339–385.

[27] M. Bodirsky, P. Jonsson, B. Martin, and A. Mottet, *Classification transfer for qualitative reasoning problems*, in Proc. 27th International Joint Conference on Artificial Intelligence (IJCAI-2018), 2018, pp. 1256–1262.

[28] M. Bodirsky, P. Jonsson, and T. V. Pham, *The complexity of phylogeny constraint satisfaction problems*, ACM Transactions on Computational Logic, 18 (2017), pp. 23:1–23:42.

[29] M. Bodirsky and J. Kára, *The complexity of temporal constraint satisfaction problems*, in Proceedings of

the Annual Symposium on Theory of Computing (STOC), C. Dwork, ed., ACM, May 2008, pp. 29–38.
[30] M. Bodirsky and J. Kára, *The complexity of temporal constraint satisfaction problems*, Journal of the ACM, 57 (2009), pp. 1–41.
[31] M. Bodirsky and J. Kára, *The complexity of temporal constraint satisfaction problems*, J. ACM, 57 (2010), pp. 9:1–9:41, https://doi.org/10.1145/1667053.1667058, https://doi.org/10.1145/1667053.1667058.
[32] M. Bodirsky and J. Kára, *A fast algorithm and Datalog inexpressibility for temporal reasoning*, ACM Transactions on Computational Logic, 11 (2010).
[33] M. Bodirsky and S. Knäuer, *Network satisfaction for symmetric relation algebras with a flexible atom*, in Proc. 35th AAAI Conference on Artificial Intelligence (AAAI-2021), 2021, pp. 6218–6226.
[34] M. Bodirsky, F. R. Madelaine, and A. Mottet, *A proof of the algebraic tractability conjecture for monotone monadic SNP*, SIAM Journal on Computing, 50 (2021), pp. 1359–1409.
[35] M. Bodirsky, B. Martin, M. Pinsker, and A. Pongrácz, *Constraint satisfaction problems for reducts of homogeneous graphs*, SIAM Journal on Computing, 48 (2019), pp. 1224–1264.
[36] M. Bodirsky and A. Mottet, *A dichotomy for first-order reducts of unary structures*, Logical Methods in Computer Science, 14 (2018), https://doi.org/10.23638/LMCS-14(2:13)2018.
[37] M. Bodirsky and J. Nešetřil, *Constraint satisfaction with countable homogeneous templates*, in Proceedings of Computer Science Logic (CSL), Vienna, 2003, pp. 44–57.
[38] M. Bodirsky and J. Nešetřil, *Constraint satisfaction with countable homogeneous templates*, Journal of Logic and Computation, 16 (2006), pp. 359–373.
[39] M. Bodirsky, W. Pakusa, and J. Rydval, *Temporal constraint satisfaction problems in fixed-point logic*, in Proc. 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS-2020), 2020, pp. 237–251.
[40] M. Bodirsky, W. Pakusa, and J. Rydval, *On the descriptive complexity of temporal constraint satisfaction problems*, Journal of the ACM, (2022). A conference version of this article appear under the title "Temporal Constraint Satisfaction Problems in Fixed-Point Logic" in LICS'20.
[41] M. Bodirsky and M. Pinsker, *Schaefer's theorem for graphs*, Journal of the ACM, 62 (2015), pp. 19:1–19:52.
[42] M. Bodirsky and M. Pinsker, *Topological Birkhoff*, Transactions of the American Mathematical Society, 367 (2015), pp. 2527–2549.
[43] M. Bodirsky and M. Pinsker, *Canonical functions: a proof via topological dynamics*, Contributions to Discrete Mathematics, 16 (2021), pp. 36–45.
[44] M. Bodirsky, M. Pinsker, and A. Pongrácz, *Projective clone homomorphisms*, Journal of Symbolic Logic, 86 (2021), pp. 148–161.
[45] M. Bodirsky, M. Pinsker, and T. Tsankov, *Decidability of definability*, Journal of Symbolic Logic, 78 (2013), pp. 1036–1054.
[46] A. A. Bulatov, *A dichotomy theorem for nonuniform CSPs*, in Proc. 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS-2017), 2017, pp. 319–330.
[47] A. A. Bulatov and E. S. Skvortsov, *Amalgams of constraint satisfaction problems*, in Proc. 18th International Joint Conference on Artificial Intelligence (IJCAI-2003), 2003, pp. 197–202.
[48] P. J. Cameron, *Transitivity of permutation groups on unordered sets*, Mathematische Zeitschrift, 148 (1976), pp. 127–139.
[49] S.-K. Chang and E. Jungert, *Symbolic Projection for Image Information Retrieval and Spatial Reasoning*, Academic Press, New York, 1996.
[50] A. Chernikov, *Lecture Notes on Stability Theory*, AMS Open Math Notes, 2019.
[51] A. G. Cohn, J. Renz, and M. Sridhar, *Thinking inside the box: A comprehensive spatial representation for video analysis*, in Proc. 13th International Conference on Principles of Knowledge Representation and Reasoning (KR-2012), 2012.
[52] T. L. Dean and M. S. Boddy, *Reasoning about partially ordered events*, Artificial Intelligence, 36 (1988), pp. 375–399.
[53] P. Denis and P. Muller, *Predicting globally-coherent temporal structures from texts via endpoint inference and graph decomposition*, in Proc. 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011), 2011, pp. 1788–1793.
[54] I. Düntsch, *Relation algebras and their application in temporal and spatial reasoning*, Artificial Intelligence Review, 23 (2005), pp. 315–357.
[55] S. Durhan and G. Sciavicco, *Allen-like theory of time for tree-like structures*, Information and Computation, 259 (2018), pp. 375–389.
[56] F. Dylla, J. H. Lee, T. Mossakowski, T. Schneider, A. van Delden, J. van de Ven, and D. Wolter, *A survey of qualitative spatial and temporal calculi: Algebraic and computational properties*, ACM Computing Surveys, 50 (2017), pp. 7:1–7:39.
[57] E. A. Emerson and J. Y. Halpern, *"Sometimes" and "not never" revisited: on branching versus linear*

time temporal logic, Journal of the ACM, 33 (1986), pp. 151–178.

[58] T. Feder and M. Y. Vardi, *The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory*, SIAM Journal on Computing, 28 (1999), pp. 57–104.

[59] A. U. Frank, *Qualitative spatial reasoning: Cardinal directions as an example*, Int. J. Geogr. Inf. Sci., 10 (1996), pp. 269–290, https://doi.org/10.1080/02693799608902079, https://doi.org/10.1080/02693799608902079.

[60] M. Gavanelli, A. Passantino, and G. Sciavicco, *Deciding the consistency of branching time interval networks*, in Proc. 25th International Symposium on Temporal Representation and Reasoning (TIME-2018), 2018, pp. 12:1–12:15.

[61] M. C. Golumbic and R. Shamir, *Complexity and algorithms for reasoning about time: a graph-theoretic approach*, Journal of the ACM, 40 (1993), pp. 1108 – 1133.

[62] J. Greiner, *Complexity of Constraint Satisfaction Problems for Unions of Theories*, PhD thesis, Technischen Universität Dresden, 2021.

[63] H. Guesgen, *Spatial reasoning based on Allen's temporal logic*. Report ICSI TR 89-049, International Computer Science Institute, 1989.

[64] R. Hirsch, *Expressive power and complexity in algebraic logic*, Journal of Logic and Computation, 7 (1997), pp. 309 – 351.

[65] W. Hodges, *A shorter model theory*, Cambridge University Press, Cambridge, 1997.

[66] P. G. Jeavons, *On the algebraic structure of combinatorial problems*, Theoretical Computer Science, 200 (1998), pp. 185–204.

[67] P. Jonsson and C. Bäckström, *A unifying approach to temporal constraint reasoning*, Artificial Intelligence, 102 (1998), pp. 143–155.

[68] A. Kechris, V. Pestov, and S. Todorčević, *Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups*, Geometric and Functional Analysis, 15 (2005), pp. 106–189.

[69] P. G. Kolaitis and M. Y. Vardi, *Conjunctive-query containment and constraint satisfaction*, in Proc. 17th Symposium on Principles of Database Systems (PODS-1998), 1998, pp. 205–213.

[70] M. Kompatscher and T. V. Pham, *A complexity dichotomy for poset constraint satisfaction*, IfCoLog Journal of Logics and their Applications, 5 (2018), pp. 1663–1696.

[71] M. Koubarakis, *Tractable disjunctions of linear constraints: Basic results and applications to temporal reasoning*, Theoretical Computer Science, 266 (2001), pp. 311–339.

[72] A. A. Krokhin, P. Jeavons, and P. Jonsson, *Reasoning about temporal relations: The tractable subalgebras of Allen's interval algebra*, Journal of the ACM, 50 (2003), pp. 591–640.

[73] A. H. Lachlan, *$\aleph_0$-categorical tree-decomposable structures*, Journal of Symbolic Logic, 57 (1992), pp. 501–514, https://doi.org/10.2307/2275284, https://doi.org/10.2307/2275284.

[74] L. Lamport, *The mutual exclusion problem: part I—a theory of interprocess communication*, Journal of the ACM, 33 (1986), pp. 313–326.

[75] G. Ligozat, *"Corner" relations in Allen's algebra*, Constraints, 3 (1998), pp. 165–177.

[76] G. Ligozat, *Reasoning about cardinal directions*, Journal of Visual Languages & Computing, 9 (1998), pp. 23 – 44, https://doi.org/https://doi.org/10.1006/jvlc.1997.9999, http://www.sciencedirect.com/science/article/pii/S1045926X97999997.

[77] G. Ligozat, *When tables tell it all: Qualitative spatial and temporal reasoning based on linear orderings*, in Proc. 5th International Conference on Spatial Information Theory (COSIT-2001), 2001, pp. 60–75.

[78] F. Madelaine and I. A. Stewart, *Constraint satisfaction, logic and forbidden patterns*, SIAM Journal on Computing, 37 (2007), pp. 132–163.

[79] M. Maróti and R. McKenzie, *Existence theorems for weakly symmetric operations*, Algebra Universalis, 59 (2008).

[80] R. H. Möhring, M. Skutella, and F. Stork, *Scheduling with and/or precedence constraints*, SIAM Journal on Computing, 33 (2004), pp. 393–415.

[81] A. Mottet, *A maximally tractable fragment of temporal reasoning plus successor*. MPRI Masters Dissertation, 2014.

[82] A. Mottet, T. Nagy, M. Pinsker, and M. Wrona, *Smooth approximations and relational width collapses*, in Proc. 48th International Colloquium on Automata, Languages, and Programming (ICALP-2021), 2021, pp. 138:1–138:20.

[83] A. Mottet and M. Pinsker, *Cores over Ramsey structures*, Journal of Symbolic Logic, 86 (2021), pp. 352–361.

[84] A. Mottet and M. Pinsker, *Smooth approximations and CSPs over finitely bounded homogeneous structures*, in Proc. 37th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS-2022), 2022, pp. 36:1–36:13.

[85] L. Mudrová and N. Hawes, *Task scheduling for mobile robots using interval algebra*, in Proc. 2015 IEEE International Conference on Robotics and Automation (ICRA-2015), 2015, pp. 383–388.

[86] A. Mukerjee and G. Joe, *A qualitative model for space*, in Proc. 8th National Conference on Artificial Intelligence (AAAI-1990), 1990, pp. 721–727.

[87] B. Nebel and H.-J. Bürckert, *Reasoning about temporal relations: A maximal tractable subclass of Allen's interval algebra*, Journal of the ACM, 42 (1995), pp. 43–66.

[88] J. Opatrný, *Total ordering problem*, SIAM Journal on Computing, 8 (1979), pp. 111–114.

[89] R. N. Pelavin and J. F. Allen, *A model for concurrent actions having temporal extent*, in Proc. 6th National Conference on Artificial Intelligence (AAAI-1987), 1987, pp. 246–250.

[90] M. Ragni and S. Wölfl, *Branching Allen*, in Proc. Spatial Cognition IV: Reasoning, Action, Interaction, International Conference Spatial Cognition 2004, vol. 3343 of Lecture Notes in Computer Science, 2004, pp. 323–343.

[91] D. A. Randell, Z. Cui, and A. G. Cohn, *A spatial logic based on regions and connection*, in Proc. 3rd Conference on Principles on Knowledge Representation and Reasoning (KR-1992), 1992, pp. 165–176.

[92] F. S. Regateiro, J. Bento, and J. Dias, *Floor plan design using block algebra and constraint satisfaction*, Advanced Engineering Informatics, 26 (2012), pp. 361–382.

[93] S. Shelah, *Stability, the f.c.p. and superstability; model theoretic properties of formulas in first order theory*, Annals of Mathematical Logic, 3 (1971), pp. 271–362.

[94] P. Simon, *A Guide to NIP Theories*, Lecture Notes in Logic, Cambridge University Press, 2015.

[95] F. Song and R. Cohen, *The interpretation of temporal relations in narrative*, in Proc. 7th National Conference on Artificial Intelligence (AAAI-1988), 1988, pp. 745–750.

[96] M. Westphal and S. Wölfl, *Qualitative CSP, finite CSP, and SAT: comparing methods for qualitative constraint-based reasoning*, in Proc. 21th International Joint Conference on Artificial Intelligence (IJCAI-2009), 2009, pp. 628–633.

[97] F. Zapata, V. Kreinovich, C. A. Joslyn, and E. Hogan, *Orders on intervals over partially ordered sets: extending Allen's algebra and interval graph results*, Soft Computing, 17 (2013), pp. 1379–1391.

[98] P. Zhang and J. Renz, *Qualitative spatial representation and reasoning in angry birds: The extended rectangle algebra*, in Proc. 14th International Conference on Principles of Knowledge Representation and Reasoning (KR-2014), 2014.

[99] D. Zhuk, *A proof of the CSP dichotomy conjecture*, Journal of the ACM, 67 (2020), pp. 30:1–30:78, https://doi.org/10.1145/3402029, https://doi.org/10.1145/3402029.

[100] D. N. Zhuk, *A proof of CSP dichotomy conjecture*, in Proc. 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS-2017), 2017, pp. 331–342.