

A tetrachotomy for positive first-order logic without equality

Florent Madelaine

Clermont Université, Université d'Auvergne,
LIMOS, BP 10448, F-63000 Clermont-Ferrand
Email: florent.madelaine@u-clermont1.fr
CNRS, UMR 6158, LIMOS, F-63173 Aubière.

Barnaby Martin

School of Engineering and Computing Sciences,
Durham University, South Road, Durham DH1 3LE, U.K.
Email: barnabymartin@gmail.com
Supported by EPSRC grant EP/G020604/1.

Abstract—We classify completely the complexity of evaluating positive equality-free sentences of first-order logic over a fixed, finite structure \mathcal{D} . This problem may be seen as a natural generalisation of the quantified constraint satisfaction problem QCSP(\mathcal{D}). We obtain a tetrachotomy for arbitrary finite structures: each problem is either in L, is NP-complete, is co-NP-complete or is Pspace-complete. Moreover, its complexity is characterised algebraically in terms of the presence or absence of specific surjective hyper-endomorphisms; and, logically, in terms of relativisation properties with respect to positive equality-free sentences.

We prove that the meta-problem, to establish for a specific \mathcal{D} into which of the four classes the related problem lies, is NP-hard. **Keywords:** Galois Connection, Quantified Constraint Satisfaction, Universal Algebra, Computational Complexity, Logic in Computer Science.

I. INTRODUCTION

We conclude the study of the complexity of the evaluation problem for positive equality-free first-order logic, on fixed and finite structures, begun in [1], [2]. This problem, denoted $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}), is a close relative of the *constraint satisfaction problem*, CSP(\mathcal{D}), and an even closer relative of the *quantified CSP*, QCSP(\mathcal{D}). In fact, it is noted in [1] that among a wide family of problems, the only interesting case, other than the CSP and QCSP, is the one addressed in the present paper (the other cases fail to manifest complexity-theoretic richness). The bulk of the theoretical research into CSPs concerns the so-called *dichotomy conjecture*: that the complexity of the problem of evaluating a primitive positive sentence on a fixed finite \mathcal{D} , CSP(\mathcal{D}), is either in P or is NP-complete. This was solved for structures with two-element domains in [3] and improved to encompass structures with three-element domains in [4]. The most successful approach to date, and the method used in [4], has been the so-called algebraic method, in which the problem of classification reverts to classes of functions under which the relevant relational structures are invariant. A similar algebraic approach has been successful in the study of the evaluation problem for positive Horn sentences, QCSP(\mathcal{D}), and, while no formal trichotomy has there been conjectured, the only known attainable complexities are P, NP-complete and Pspace-complete [5], [6].

The problem $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) may be seen as a generalisation of CSP(\mathcal{D}) (which becomes in our notation $\{\exists, \wedge, =$

$\}$ -FO(\mathcal{D})) and QCSP(\mathcal{D}) ($\{\exists, \forall, \wedge, =\}$ -FO(\mathcal{D})), despite the forbidding of equality – as in instances of CSP and QCSP, equalities may be propagated out in all but trivial cases. In contrast, the forbidding of equality from $\{\exists, \forall, \wedge, \vee, =\}$ -FO(\mathcal{D}) becomes very significant, as these problems with equality have a straightforward complexity classification [7] (note that equalities may no longer be propagated out). In [1] and [2], the complexity of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is studied through an analogous algebraic method to that used for CSP(\mathcal{D}). These papers culminate in a full classification – a tetrachotomy – as \mathcal{D} ranges over structures with four-element domains. Specifically, the problems $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) are either in L, are NP-complete, are co-NP-complete or are Pspace-complete. In [2], a conjecture is ventured as to where the general tetrachotomy lies for arbitrary finite structures \mathcal{D} . In this paper we prove that conjecture.

We derive our result by studying sets of surjective hyper-operations (shops) under which our relational \mathcal{D} may be invariant. These sets, always containing the identity and closed under composition and sub-shops, are known as *down-shop-monoids*. For membership of L, NP and co-NP, it is proved in [8] that it is sufficient to have certain special shops as surjective hyper-endomorphisms (shes). The converse, that it is necessary to have these special shops as shes, is more subtle, and is the major contribution of this paper.

For the CSP there is the well-known notion of the core of \mathcal{D} , that may be seen as the minimal induced substructure $\mathcal{D}' \subseteq \mathcal{D}$ such that \mathcal{D}' and \mathcal{D} agree on all primitive positive sentences. Equivalently, the domain D' of \mathcal{D}' is minimal such that any primitive positive sentence is true on \mathcal{D} iff it is true on \mathcal{D} with all (existential) quantifiers relativised to D' . Cores are minimal structures in their equivalence classes, given by the equivalence relation of satisfying the same primitive positive sentences. Cores are very robust, for instance, being unique up to isomorphism, and sitting as induced substructures in all other structures in their equivalence class. A similar notion to core exists for the QCSP, but it is not nearly so robust (they need no longer be unique nor sit as an induced substructure in other structures in their equivalence class [9]). For the problems $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}), a notion of core returns, and it is once again robust. The *U-X-core* of \mathcal{D} consists of a minimal substructure induced by the union $U \cup X$ of two

minimal sets U and X of D such that a positive equality-free sentence is true on \mathcal{D} iff it is true on \mathcal{D} with the universal quantifiers relativised to U and the existential quantifiers relativised to X . Analysing U - X -cores gives us the necessary converse alluded to in the previous paragraph. In the Pspace-complete case, some completion of the U - X -core is either fundamentally very simple and can be classified as in a two-element domain, known from [1], or it is a generalisation of one of the four-element cases from [2]. For the NP-complete and co-NP-complete cases, some completion of the U - X -core is fundamentally very simple and can be classified as an easy generalisation of a three-element domain.

We are able therefore to give the delineation of our tetrachotomy by two equivalent means. Firstly, by the presence or absence of certain shops – so-called *A-shops* and *E-shops* – as shes. Secondly, by the existence or not of trivial sets for the relativisation of universal and existential quantifiers. Thus, $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is in L iff \mathcal{D} has both an A-shop and an E-shop for shes, iff there exist singleton sets U and X such that a sentence of positive equality-free logic is true on \mathcal{D} exactly when it is true on \mathcal{D} with the universal quantifiers and existential quantifiers relativised to U and X , respectively. Otherwise, and in a similar vein, $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is NP-complete (resp., co-NP-complete) if it has an A-shop (resp., E-shop) for a she, iff there exists a singleton set U (resp., X) such that a sentence of positive equality-free logic is true on \mathcal{D} exactly when it is true on \mathcal{D} with the universal quantifiers relativised to U (resp., the existential quantifiers relativised to X). In all remaining cases, $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is Pspace-complete, and \mathcal{D} has neither an A-shop nor an E-shop as a she, and there are no trivial sets U nor X affording the required relativisation properties.

Finally, we investigate the complexity of the meta-problem: given a finite structure \mathcal{D} , what is the complexity of evaluating positive equality-free sentences of FO over \mathcal{D} ? We establish that the meta-problem is NP-hard, even for a fixed and finite signature.

The paper is organised as follows. After the preliminaries of Section II, we introduce the central notion of U - X -cores, and relativisation, in Section III. In Section IV, we use cores to derive the outstanding NP-, co-NP- and Pspace-hardness proofs. In Section V, we examine the complexity of the meta-problem and in Section VI we make some final observations. For reasons of space, several proofs are deferred to the appendix.

II. PRELIMINARIES

Unless otherwise stated, we shall work with finite relational structures that share the same finite relational signature σ . Let \mathcal{D} be such a structure. We will denote its domain by D . We denote the size of such a set D by $|D|$. The *complement* $\overline{\mathcal{D}}$ of a structure \mathcal{D} consists of relations that are exactly the set-theoretic complements of those in \mathcal{D} . I.e., for an a -ary R , $R^{\overline{\mathcal{D}}} := D^a \setminus R^{\mathcal{D}}$. For graphs this leads to a slightly non-standard notion of complement, as it includes self-loops.

We denote positive equality-free first-order logic by $\{\exists, \forall, \wedge, \vee\}$ -FO. We will occasionally refer to other logics in this notation, for example *primitive positive* logic is $\{\exists, \wedge, =\}$ -FO. For a fixed and finite \mathcal{D} , the problem $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) takes as input a sentence φ of $\{\exists, \forall, \wedge, \vee\}$ -FO and asks whether $\mathcal{D} \models \varphi$. We will occasionally refer to related problems in this notation, for example $\{\exists, \wedge, =\}$ -FO(\mathcal{D}) – better known as CSP(\mathcal{D}) – takes as input a sentence φ of primitive positive logic, and asks whether $\mathcal{D} \models \varphi$.

Let D be a finite set. A *hyper-operation* is a function from D to its powerset. When we wish to stress that an element may be sent to \emptyset , we speak of a *partial hyper-operation*, (phop). For a subset S of D , we will define its image $f(S)$ under the hyper-operation f as $\bigcup_{s \in S} f(s)$. Let f and g be two hyper-operations. The hyper-operation $g \circ f$ is defined naturally as $g \circ f(x) := g(f(x))$ (recall that $f(x)$ is a set). A *surjective (total) hyper-operation* (shop) f over D satisfies: $f(D) = D$ (surjectivity) and $\forall z \in D, f(z) \neq \emptyset$ (totality). The inverse of a shop f , denoted by f^{-1} , is the hyper-operation defined for any y in D as $f^{-1}(y) := \{x \in D \mid y \in f(x)\}$.

We call an element of $f^{-1}(y)$ an *antecedent* of y under f . Observing that the totality of f may be rephrased as $f^{-1}(D) = D$, it is now clear that the inverse of a shop is a shop. It is also straightforward to verify that if f and g are two shops then $g \circ f$ is also a shop.

Observing that shop composition is associative and that the identity shop (which sends an element x of D to the singleton $\{x\}$) is the identity with respect to composition, we may consider the monoid generated by a set of shops. A shop f is a *sub-shop* of a shop g whenever, for every x in D , $f(x) \subseteq g(x)$. In our context, we will be interested in particular monoids which will be closed further under sub-shops, the so-called *down-shop-monoids* (DSMs).¹ We denote by $\langle F \rangle_{DSM}$ the DSM generated by a set F of shops.

Let U and X be subsets of D . We say that a shop f is *U-surjective* if $f(U) = D$. In the special case where the set U is a singleton, we speak of an *A-shop*. We say that a shop g is *X-total* if $f^{-1}(X) = D$. In the special case where the set X is a singleton, we speak of an *E-shop*. Note that these notions are dual to one another, that is the inverse of a U -surjective shop is an X -total shop with $X = U$ and vice versa.

Somewhat abusing terminology, we will drop the word surjective and by U - or U' -shop we will mean a U - or U' -surjective shop. Similarly, we will speak of an X - or X' -shop in the total (existential) case and of a U - X -shop in the case of a shop that is both U -surjective and X -total.

Suitable compositions of U -shops and X -shops preserve these properties.

Lemma 1. *Let f and g be two shops.*

- 1) *If f is a U -shop then $g \circ f$ is a U -shop.*
- 2) *If g is a X -shop then $g \circ f$ is a X -shop.*

¹The “down” comes from *down-closure*, here under sub-shops; a nomenclature inherited from [10].

- 3) If both f is a U -shop and g is a X -shop then $g \circ f$ is a U - X -shop.
- 4) If both f and g are U - X -shops then $g \circ f$ is a U - X -shop.
- 5) The iterate of a U - X -shop is a U - X -shop.

Proof: We prove (1). Since $f(U) = D$, we have $g(f(U)) = g(D)$. By surjectivity of g , we know that $g(D) = D$. It follows that $g(f(U)) = D$ and we are done. (2) is dual to (1), and (3) follows directly from (1) and (2). (4) is a restriction of (3) and is only stated here as we shall use it often. (5) follows by induction on the order of iteration using (4). \blacksquare

A *homomorphism* (resp. *strong homomorphism*) from a structure \mathcal{D} to a structure \mathcal{E} is a function $h : D \rightarrow E$ that preserves (resp. strongly preserves) the relations of \mathcal{D} , i.e. for all a_i -ary relations R_i , and for all $x_1, \dots, x_{a_i} \in D$, $R_i(x_1, \dots, x_{a_i}) \in \mathcal{D}$ implies $R_i(h(x_1), \dots, h(x_{a_i})) \in \mathcal{E}$ (resp. $R_i(x_1, \dots, x_{a_i}) \in \mathcal{D}$ iff $R_i(h(x_1), \dots, h(x_{a_i})) \in \mathcal{E}$). \mathcal{D} and \mathcal{E} are *homomorphically equivalent* if there are homomorphisms both from \mathcal{D} to \mathcal{E} and from \mathcal{E} to \mathcal{D} .

A *surjective hyper-endomorphism* (she) of \mathcal{D} is a shop f on D that satisfies, for all relations R of \mathcal{D} ,

- if $\mathcal{D} \models R(x_1, \dots, x_i)$ then, for all $y_1 \in f(x_1), \dots, y_i \in f(x_i)$, $\mathcal{D} \models R(y_1, \dots, y_i)$.

We will also say that f *preserves* \mathcal{D} . More generally, for $r_1, \dots, r_k \in D$, we say f is a *she from* $(\mathcal{D}, r_1, \dots, r_k)$ to $(\mathcal{D}, r'_1, \dots, r'_k)$ if f is a she of \mathcal{B} and $r'_1 \in f(r_1), \dots, r'_k \in f(r_k)$. A she may be identified with a *surjective endomorphism* if each element is mapped to a singleton set. On finite structures surjective endomorphisms are necessarily automorphisms. We will freely talk of A-shes and E-shes of a certain structure \mathcal{D} .

For a set F of shops on the finite domain D , let $\text{Inv}(F)$ be the set of relations on D of which each f in F is a she (when these relations are viewed as a structure over D). We say that S in $\text{Inv}(F)$ is *invariant* or *preserved* by (the shops in) F . Let $\text{shE}(\mathcal{B})$ be the set of shes of \mathcal{B} . Let $\langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}}$ be the sets of relations that may be defined on \mathcal{B} in $\{\exists, \forall, \wedge, \vee\}$ -FO. Our main methodological tool to establish the complexity of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is that the Galois connection captures the complexity [1], [11]

Theorem 2 ([1], [11] **Galois connection and Complexity**). *For a finite structure \mathcal{D} and a set of shops F , the following holds,*

- $\langle \mathcal{D} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}} = \text{Inv}(\text{shE}(\mathcal{D}))$; and,
- $\langle F \rangle_{DSM} = \text{shE}(\text{Inv}(F))$.

For finite \mathcal{D} and \mathcal{D}' (s.t. $D = D'$), if $\text{shE}(\mathcal{D}) \subseteq \text{shE}(\mathcal{D}')$ then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}') \leq_L $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}).

A. Main result

The following – left as a conjecture at the end of [2], [8] – is the main result of this paper.

Theorem 3 (Tetrachotomy). *Let \mathcal{D} be any structure.*

- I. If \mathcal{D} is preserved by both an A-shop and an E-shop, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is in L.

- II. If \mathcal{D} is preserved by an A-shop but is not preserved by any E-shop, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is NP-complete.
- III. If \mathcal{D} is preserved by an E-shop but is not preserved by any A-shop, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is co-NP-complete.
- IV. If \mathcal{D} is preserved neither by an A-shop nor by an E-shop, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is Pspace-complete.

Proof: The upper bounds (membership in L, NP and co-NP) for Cases I, II and III were known from [8], but we reprove them here as Corollary 9. Theorem 18 deals with the lower bounds. NP-hardness for Case II and co-NP-hardness for Case III are proved in Subsection IV-C. Pspace-hardness for Case III is proved in Subsection IV-D. \blacksquare

III. THE NOTION OF A CORE FOR POSITIVE EQUALITY-FREE FO

Recall that a *core* of a finite structure \mathcal{D} is an induced substructure $\tilde{\mathcal{D}}$ of \mathcal{D} such that there is a homomorphism from \mathcal{D} to $\tilde{\mathcal{D}}$ (and consequently, they are homomorphically equivalent) and every endomorphism of $\tilde{\mathcal{D}}$ is an automorphism. The core of a structure is unique up to isomorphism. The notion of core is ubiquitous in the study of CSPs. The property of core can be rephrased in the logical context as the minimal $X = \tilde{D} \subseteq D$ such that a primitive positive sentence φ is true on \mathcal{D} iff it is true on \mathcal{D} with the (existential) quantifiers relativised to $X = \tilde{D}$. Let us say in this case that \mathcal{D} has *X-relativisation* with respect to primitive positive logic $\{\exists, \wedge, =\}$ -FO.

Thus, the notion of a core can be recast in the context of $\{\exists, \wedge, =\}$ -FO in a number of equivalent ways, as a minimal induced substructure $\tilde{\mathcal{D}}$ of \mathcal{D} ,

- that satisfies the same $\{\exists, \wedge, =\}$ -FO sentences;
- that is induced by minimal $X \subseteq D$ such that \mathcal{D} has X -relativisation w.r.t. $\{\exists, \wedge, =\}$ -FO; or,
- that is induced by minimal $X \subseteq D$ such that \mathcal{D} has an endomorphism with image X .

We are looking for a useful characterisation of the analogous concept for our positive equality-free logic, $\{\exists, \forall, \wedge, \vee\}$ -FO. As we now have both quantifiers, it is natural to specify two sets U and X , one for each quantifier. As we shall see shortly, there are two equivalent ways of defining a U - X -core – one is logical, the other algebraic – as a minimal substructure $\tilde{\mathcal{D}}$ of \mathcal{D} , induced by minimal $U, X \subseteq D$ such that:

- \mathcal{D} has U - X -relativisation w.r.t. $\{\exists, \forall, \wedge, \vee\}$ -FO; or,
- \mathcal{D} has a U - X -she.

We will show that the sets U and X are unique up to isomorphism and that within a minimal induced substructure $\tilde{\mathcal{D}}$, the sets U and X are uniquely determined. This will reconcile our definition of a U - X -core with the following natural definition, in which U and X are not explicit:

- as a minimal induced substructure $\tilde{\mathcal{D}}$ of \mathcal{D} that satisfies the same $\{\exists, \forall, \wedge, \vee\}$ -FO sentences.

In order to prove the equivalence between the algebraic and the logical definitions we need some technical results about U - X -shops. Since it is U - X -relativisation that motivated us

to introduce U - X -shops in the first place, we will delay until later the proof of these results.

A. Relativisation

Given a formula φ , we denote by $\varphi_{[\forall u/\forall u \in U, \exists x/\exists x \in X]}$ the formula obtained from φ by relativising simultaneously every universal quantifier to U and every existential quantifier to X . When we only relativise universal quantifiers to U , we write $\varphi_{[\forall u/\forall u \in U]}$, and when we only relativise existential quantifiers to X , we write $\varphi_{[\exists x/\exists x \in X]}$. We will say that a shop g over D may act as the identity over $Z \subseteq D$ iff for every $z \in Z$, we have $g(z) \ni z$. Note that this property is preserved by shop composition.

Definition 4. Let \mathcal{D} be a finite structure over a set D , and U, X be two subsets of D . We say that \mathcal{D} has U - X -relativisation if, for all formulae $\varphi(\mathbf{w})$ in $\{\exists, \forall, \wedge, \vee\}$ -FO whose free variables \mathbf{w} may range over $U \cup X$, the following are equivalent

- (i) $\mathcal{D} \models \varphi(\mathbf{w})$
- (ii) $\mathcal{D} \models \varphi(\mathbf{w})_{[\forall u/\forall u \in U]}$
- (iii) $\mathcal{D} \models \varphi(\mathbf{w})_{[\exists x/\exists x \in X]}$
- (iv) $\mathcal{D} \models \varphi(\mathbf{w})_{[\forall u/\forall u \in U, \exists x/\exists x \in X]}$

Lemma 5. Let \mathcal{D} be a finite structure over a set D , and U, X be two subsets of D . If \mathcal{D} has a U - X -she that may act as the identity over $U \cup X$, then \mathcal{D} has U - X -relativisation.

Proof: Let g be a U - X -she of \mathcal{D} such that for every z in $U \cup X$, we have $g(z) \ni z$. This means that when taking the image or a preimage under this she g , we may fix at will elements in $U \cup X$.

We have $(iii) \Rightarrow (i) \Rightarrow (ii)$ and $(iii) \Rightarrow (iv) \Rightarrow (ii)$ trivially. It suffices to prove that $(ii) \Rightarrow (i)$ and $(i) \Rightarrow (iii)$ to complete the proof. To do so, we will consider the well known Hintikka game corresponding to Case (i), called the *unrelativised game* hereafter; and, the relativised Hintikka games corresponding to the relativised formulae from Cases (ii), (iii) and (iv) (the relativised game considered being clear from context). We will show how to transfer existence of winning strategies from one game to another. We will not introduce formally the notion of a winning strategy, as the heavy notation would not clarify the proof exposition.

$(ii) \Rightarrow (i)$. If \forall plays a variable v_\forall on some value d , such that $g(u') \ni d$ for some u' in U (since g is a U -shop), then \exists will consider that he has actually played on u' in the relativised game, for which \exists has a winning strategy, and can answer some existential variable v_\exists with some $e' \in D$. In the unrelativised game, \exists now answers with some $e \in D$ such that $g(e') \ni e$ (since g is total). Note that applying the shop g to \exists 's relativised winning strategy, sending u' to d and e' to e , leaving free variables fixed, and preserving positive information, proves that \exists is indeed following a winning strategy.

$(i) \Rightarrow (iii)$. If \forall plays a variable v_\forall on some value d , such that $g(d') \ni d$ for some d' in D (since g is surjective), then \exists has a winning strategy in the unrelativised case and can

answer some existential variable v_\exists with some $e' \in D$. In the relativised game, \exists now answers to \forall 's move of d for v_\forall with some $x \in X$ such that $g(e') \ni x$ (since g is an X -shop). Note that applying the shop g to \exists 's relativised winning strategy, sending d' to d and e' to x , and leaving free variables fixed, proves that \exists is indeed following a winning strategy. \blacksquare

We will prove the converse shortly. We will use the formula $\Theta_{\mathbf{r}}$ from the following lemma, which was used to establish one half of the Galois Connection in [1].

Lemma 6 ([1] Canonical $\{\exists, \forall, \wedge, \vee\}$ -FO-query). Let $\mathbf{r} := (r_1, \dots, r_k)$ be a k -tuple of elements of \mathcal{D} . There exists a formula $\Theta_{\mathbf{r}}(u_1, \dots, u_k) \in \{\exists, \forall, \wedge, \vee\}$ -FO s.t. $(\mathcal{D}, r'_1, \dots, r'_k) \models \Theta_{\mathbf{r}}(u_1, \dots, u_k)$ iff there is a she from $(\mathcal{D}, r_1, \dots, r_k)$ to $(\mathcal{D}, r'_1, \dots, r'_k)$.

We will now explain what this formula Θ involves (for more see [1]). Let $\mathbf{r} \in D^k$, $\mathbf{s} := (d_1, \dots, d_{|D|})$ be an enumeration of D and $\mathbf{t} \in D^{|D|}$. $\Phi_{\mathcal{D}(\mathbf{r}, \mathbf{s})}(u_1, \dots, u_k, v_1, \dots, v_{|D|})$ is a conjunction of the positive facts of (\mathbf{r}, \mathbf{s}) , where the variables (\mathbf{u}, \mathbf{v}) correspond to the elements (\mathbf{r}, \mathbf{s}) . Similarly, $\Phi_{\mathcal{D}(\mathbf{r}, \mathbf{s}, \mathbf{t})}(u_1, \dots, u_k, v_1, \dots, v_{|D|}, w_1, \dots, w_{|D|})$ is the conjunction of the positive facts of $(\mathbf{r}, \mathbf{s}, \mathbf{t})$, where the variables $(\mathbf{u}, \mathbf{v}, \mathbf{w})$ correspond to the elements $(\mathbf{r}, \mathbf{s}, \mathbf{t})$. Set

$$\Theta_{\mathbf{r}}(\mathbf{u}) := \exists \mathbf{v} \Phi_{\mathcal{D}(\mathbf{r}, \mathbf{s})}(\mathbf{u}, \mathbf{v}) \wedge \forall \mathbf{w} \bigvee_{\mathbf{t} \in D^{|D|}} \Phi_{\mathcal{D}(\mathbf{r}, \mathbf{s}, \mathbf{t})}(\mathbf{u}, \mathbf{v}, \mathbf{w}).$$

We have already pointed out that U -shop and X -shop are dual to one another. It is a simple exercise to show that the complement $\overline{\mathcal{D}}$ of a structure \mathcal{D} is preserved precisely by the inverses of the shops that preserve \mathcal{D} [11]. Thus the above lemma holds in reverse for a formula $\Xi_{\mathbf{r}}$ that is built in exactly the same manner as $\Theta_{\mathbf{r}}$ but by using conjunctive *negated* queries Ψ involving negative facts instead of Φ and its positive facts (in other words, the negative fragment of equality-free FO is characterised by the same lattice of DSMs).

Lemma 7. Let \mathcal{D} be a finite structure over a set D , and U, X be two subsets of D . If \mathcal{D} has U - X -relativisation then \mathcal{D} has a U - X -she that may act as the identity over $U \cup X$.

Proof: We will use $\Theta_{\mathbf{r}}$ and $\Xi_{\mathbf{r}}$ to establish that \mathcal{D} has an X -she g and a U -she f , that may both act as the identity over $U \cup X$. Then we may compose these to obtain a U - X -she $g \circ f$ that may act as the identity over $U \cup X$. Let $\mathbf{r} = \mathbf{r}'$ be an enumeration of $U \cup X$ (we retain the dash on \mathbf{r}' for pedagogical reasons).

Since the identity is a she from $(\mathcal{D}, \mathbf{r})$ to $(\mathcal{D}, \mathbf{r}')$, it follows by Lemma 6 that $\mathcal{D} \models \Theta_{\mathbf{r}}(\mathbf{r}')$. By assumption, we may equivalently relativise only the existential quantifiers to X ($(i) \Rightarrow (iii)$): $\mathcal{D} \models \Theta_{\mathbf{r}}(\mathbf{r}')_{[\exists x/\exists x \in X]}$. Let \mathbf{s}' be the witnesses in X for the variables \mathbf{v} . Finally we set $\mathbf{w} := \mathbf{s} = (d_1, \dots, d_{|D|})$ and obtain a witness \mathbf{t}' to the disjunction. $\mathcal{D} \models \Phi_{\mathcal{D}(\mathbf{r}, \mathbf{s})}(\mathbf{r}', \mathbf{s}') \wedge \Phi_{\mathcal{D}(\mathbf{r}, \mathbf{s}, \mathbf{t}')}(\mathbf{r}', \mathbf{s}', \mathbf{s})$. Consider the following shops on \mathcal{D} .

- 1) $g_{\mathbf{r}}$ given by $g_{\mathbf{r}}(r_i) = \{r'_i\}$, for $1 \leq i \leq |U \cup X|$.
($U \cup X$ -identity.)
- 2) $g_{\mathbf{s}}$ given by $g_{\mathbf{s}}(d_i) = \{s'_i\}$, for $1 \leq i \leq |D|$. (X -totality.)

- 3) g_t given by $d_i \in g_t(d_j)$ iff $t'_i = d_j$, for $1 \leq i, j \leq |D|$. (surjectivity.)

Let $g := g_r \cup g_s \cup g_t$. By construction, g is an X -shop that may act as the identity over $U \cup X$ by g_r . Its surjectivity is guaranteed by g_t and its “ X -totality” is guaranteed by g_s . That g is a she of \mathcal{D} follows from the right-hand conjunct $\Phi_{\mathcal{D}(r,s,t)}$.

Now recall that Ξ_r is

$$\exists \mathbf{v} \Psi_{\mathcal{D}(r,s)}(\mathbf{u}, \mathbf{v}) \wedge \forall \mathbf{w} \bigvee_{\mathbf{t} \in D^{|D|}} \Psi_{\mathcal{D}(r,s,t)}(\mathbf{u}, \mathbf{v}, \mathbf{w}).$$

Since the identity is a she from $(\mathcal{D}, \mathbf{r}')$ to $(\mathcal{D}, \mathbf{r})$, it follows that $\mathcal{D} \models \Xi_r(\mathbf{r}')$. Now, this is equivalent to

$$\mathcal{D} \not\models \forall \mathbf{v} \neg \Psi_{\mathcal{D}(r,s)}(\mathbf{r}', \mathbf{v}) \vee \exists \mathbf{w} \bigwedge_{\mathbf{t} \in D^{|D|}} \neg \Psi_{\mathcal{D}(r,s,t)}(\mathbf{r}', \mathbf{v}, \mathbf{w}).$$

This is now a sentence of $\{\exists, \forall, \wedge, \vee\}$ -FO and by the U - X -relativisation assumption, we may relativise only the universal quantifiers to U (contrapositive of (ii) \Rightarrow (i)). Rewriting, we obtain

$$\mathcal{D} \models \exists \mathbf{v} \in U \Psi_{\mathcal{D}(r,s)}(\mathbf{r}', \mathbf{v}) \wedge \forall \mathbf{w} \bigvee_{\mathbf{t} \in D^{|D|}} \Psi_{\mathcal{D}(r,s,t)}(\mathbf{r}', \mathbf{v}, \mathbf{w}).$$

Now we obtain witnesses \mathbf{s}' in U and set $\mathbf{w} := \mathbf{s} = (d_1, \dots, d_{|D|})$ to obtain a witness \mathbf{t}' to the disjunction.

$$\mathcal{D} \models \Psi_{\mathcal{D}(r,s)}(\mathbf{r}', \mathbf{s}') \wedge \Psi_{\mathcal{D}(r,s,t')}(\mathbf{r}', \mathbf{s}', \mathbf{s}).$$

Consider the following phops on \mathcal{D} .

- 1) f_r given by $f_r(r'_i) = \{r_i\}$, for $1 \leq i \leq |U \cup X|$. ($U \cup X$ -identity.)
- 2) f_s given by $f_s(s'_i) \ni d$, for $1 \leq i \leq |D|$. (U -surjectivity.)
- 3) f_t given by $f_t(d_i) = \{d_j\}$ iff $t'_i = d_j$, for $1 \leq i, j \leq |D|$. (totality.)

Let $f := f_r \cup f_s \cup f_t$. By construction, f is an X -shop that may act as the identity over $U \cup X$ by f_r . Its “ U -surjectivity” is guaranteed by f_t and its totality is guaranteed by f_s . That f is a she of \mathcal{D} follows from the right-hand conjunct $\Psi_{\mathcal{D}(r,s,t')}$. ■

Together, the two previous lemmata establish an algebraic characterisation of relativisation.

Theorem 8. *Let \mathcal{D} be a finite structure over a set D , and U, X be two subsets of D . The structure \mathcal{D} has U - X -relativisation if, and only if, it has a U - X -shop as a she that may act as the identity over $U \cup X$.*

As a corollary, we obtain the upper bound for the complexity in Cases I, II and III of our tetrachotomy theorem (Theorem 3). We knew this already [8] but it keeps this paper self-contained.

Corollary 9 (upper bounds).

- I. If \mathcal{D} is preserved by both an A -shop and an E -shop then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is in L .
- II. If \mathcal{D} is preserved by an A -shop then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is in NP .
- III. If \mathcal{D} is preserved by an E -shop then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is in $co-NP$.

Proof: We first prove Case II. Recall that an A -shop is simply a U - X -shop with $U = \{u\}$, for some u in D , and $X \subseteq D$. As we shall prove shortly (see Proposition 16), there is a subset X' of X such that there is a U - X' -shop which may act as the identity over $U \cup X'$ as required. We may therefore replace every universal quantifier by the constant u and relativise every existential quantifier to X' by Theorem 8. This means that $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is in NP in the presence of an A -she.

Recall that an E -shop is simply a U - X -shop with $X = \{x\}$ for some x in D , and $U \subseteq D$. So Case III is dual to Case II and we finally turn to Case I. With both an A -shop and an E -shop, we have a U - X -shop with $U = \{u\}$ and $X = \{x\}$ where u and x are in D . Plainly this U - X -shop may act as the identity over $U \cup X$. We may therefore replace every universal quantifier by the constant u and every existential quantifier by the constant x , by Theorem 8. We have reduced $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) to the Boolean sentence value problem, known to be in L [12]. ■

An induced substructure \mathcal{D}' of \mathcal{D} , such that there is an endomorphism from \mathcal{D} to \mathcal{D}' which is the identity over D' , is called a *retract*; and cores are simply minimal retracts. In the context of $\{\exists, \forall, \wedge, \vee\}$ -FO, the analogous notion corresponds to a substructure induced by a she which is a U - X -shop that may act as the identity over $U \cup X$.

Corollary 10 (retraction). *Let \mathcal{D} be a finite structure that has a U - X -she that may act as the identity over $U \cup X$. Let $\tilde{\mathcal{D}}$ be the substructure of \mathcal{D} induced by $U \cup X$. For any sentence φ in $\{\exists, \forall, \wedge, \vee\}$ -FO, $\mathcal{D} \models \varphi$ if, and only if, $\tilde{\mathcal{D}} \models \varphi$.*

B. The U - X Core.

Let \mathcal{D} be a finite structure and \mathcal{M} its associated DSM; i.e. $\mathcal{M} := \text{shE}(\mathcal{D})$.

We consider all minimal subsets X of D such that there is an X -shop g in \mathcal{M} , and all minimal subsets U such that there is a U -shop f in \mathcal{M} . Such sets always exist as totality and surjectivity of shops mean that in the worst case we may choose $U = X = D$. Since \mathcal{M} is a monoid, by virtue of Lemma 1, it follows that $g \circ f$ is a U - X shop and belongs to \mathcal{M} . Thus, we may furthermore require that among all minimal sets satisfying the above, we choose a set U and a set X with $U \cap X$ maximal. Let $\tilde{\mathcal{D}}$ be the substructure of \mathcal{D} induced by $U \cup X$. We call $\tilde{\mathcal{D}}$ a U - X -core of \mathcal{D} .

Note that the sets U and X are not necessarily unique. However, as we shall see later the U - X -core is unique up to isomorphism, so we fix them arbitrarily from now on.

Remark. Assume that there is an X_1 -shop h_1 and an X_2 -shop h_2 in \mathcal{M} s.t. $|X_1| > |X_2|$. We consider images of $h_1 \circ h_2$:

$$D \xrightarrow{h_2} X_2 \subseteq h_2(D) \xrightarrow{h_1} \emptyset \subsetneq h_1(X_2) \cap X_1 \subseteq X_1 \subseteq h_1 \circ h_2(D).$$

For each element in X_2 , pick a single element x'_1 of X_1 in $h_1(X_2) \cap X_1$ s.t. $x'_1 \in h_1(x_2)$. Let X'_1 denote the set of picked elements. Since $|X_1| > |X_2|$ then $h_1 \circ h_2$ is an X'_1 -shop in \mathcal{M} with $|X'_1| \leq |X_2|$.

This means that we may look for an X -shop in \mathcal{M} where the set X is minimal w.r.t. inclusion, or equivalently, for a set with minimal size $|X|$. So, in order to find an X -shop with a minimal set $|X|$ in \mathcal{M} , we can remove one by one an element from D until we obtain a set X such that \mathcal{M} contains an X -shop, but no X' -shop for $X' \subsetneq X$. The dual argument applies to U -shops, and consequently to U - X -shops.

This further explains why minimising U and X , and then maximising their intersection, necessarily leads to a minimal $\widehat{D} := U \cup X$ also. Because, would we find $U' \cup X'$ of smaller size, we might look within U' and X' for potentially smaller sets of cardinality $|U|$ and $|X|$, thus contradicting minimality.

Lemma 11. *Let f be a shop in \mathcal{M} . For any element z in D , $f(z)$ contains at most one element of the set U , that is $|f(z) \cap U| \leq 1$.*

Proof: Assume for contradiction that there is some z and some distinct elements u_1 and u_2 of U such that $f(z) \supseteq \{u_1, u_2\}$. Let z_3, z_4, \dots be any choice of antecedents under f of the remaining elements u_3, u_4, \dots of U (recall that f is surjective). By assumption the monoid \mathcal{M} contains a U -shop g . Hence, $g \circ f$ would be a U' -shop with $U' = \{z, z_3, z_4, \dots\}$ since $f(U') \subseteq U$ and $g(U) = D$. We get a contradiction as $|U'| < |U|$. ■

Lemma 12. *Let f be a U -shop in \mathcal{M} . There exists a permutation α of U such that: for any u in U ,*

- $f(u) \cap U = \{\alpha(u)\}$; and,
- $f^{-1}(u) \cap U = \{\alpha^{-1}(u)\}$.

The dual statements concerning X -shops hold.

Lemma 13. *Let f be a shop in \mathcal{M} . for any element z in D , $f^{-1}(z)$ contains at most one element of the set X , that is $|f^{-1}(z) \cap X| \leq 1$.*

Lemma 14. *Let f be an X -shop in \mathcal{M} . There exists a permutation β of X such that: for any x in X ,*

- $f(x) \cap X = \{\beta(x)\}$; and,
- $f^{-1}(x) \cap X = \{\beta^{-1}(x)\}$.

Lemma 15. *Let f be a shop in \mathcal{M} . If f is a U - X -shop then $f(X) \cap (U \setminus X) = \emptyset$.*

Proof: Assume for contradiction that for some $x_1 \in X$ and some $u_1 \in U \setminus X$, we have $u_1 \in f(x_1)$. Since f is an X -shop, every element is an antecedent under f of some element in X , in particular every element $x_2, x_3, \dots \in X$ (different from x_1) has a unique image $x'_2, x'_3, \dots \in X$ (see Lemma 14). Some element of X , say x_i does not occur in these images. Necessarily, x_1 reaches x_i . Note that x_i can not also belong to U as otherwise, x_i and u_1 , two distinct elements of U , would be reached by x_1 , contradicting Lemma 11. Thus, we must have that x_i belongs to $X \setminus U$. Let $U' := U$ and $X' := X \setminus \{x_i\} \cup \{u_1\}$. Note that $f^2 := f \circ f$, the second iterate of f , is a U' - X' -shop with $|U'| = |U|$, $|X'| = |X|$ and $|U' \cap X'| < |U \cap X|$. This contradicts our hypothesis on U and X . ■

Proposition 16. *There is a U - X -shop h in \mathcal{M} that may act as the identity on $U \cup X$ and has the properties:*

- for any y in $U \cap X$, $h(y) \cap (U \cup X) = \{y\}$;
- for any x in $X \setminus U$, $h(x) \cap (U \cup X) = \{x\}$;
- for any u in $U \setminus X$, $h(u) \cap (U \cup X) = \{u\} \cup X_u$, where $X_u \subseteq X \setminus U$; and,
- $h(U \setminus X) \cap X = \bigcup_{u \in U \setminus X} X_u = X \setminus U$.

We do not need the following to prove the tetrachotomy.

Theorem 17. *The U - X -core is unique up to isomorphism.*

IV. PROVING HARDNESS

Our aim is to derive the following lower bounds.

- Theorem 18 (lower bounds).** II. *If \mathcal{D} is preserved by an A -shop but is not preserved by any E -shop, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is NP-hard.*
- III. *If \mathcal{D} is preserved by an E -shop but is not preserved by any A -shop, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is co-NP-hard.*
- IV. *If \mathcal{D} is preserved neither by an A -shop nor by an E -shop, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is Pspace-hard.*

It follows from Proposition 16 and Corollary 10 that the complexity of a structure \mathcal{D} is the same as the complexity of its U - X -core. Hence in this Section, we assume w.l.o.g. that $U \cup X = D$. We will say in this case that the DSM \mathcal{M} is reduced. In order to prove Theorem 18, we need to establish the following:

- II. If U is of size one and X of size at least two then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is NP-hard;
- III. If X is of size one and U of size at least two then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is co-NP-hard; and,
- IV. If both U and X have at least two elements then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is Pspace-hard.

In the following, we will describe a DSM \mathcal{M} as being (NP-, co-NP-, Pspace-)hard in the case that $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is hard for some $\mathcal{D} \in \text{Inv}(\mathcal{M})$. In order to facilitate the hardness proof, we would like to show hardness of a monoid $\widehat{\mathcal{M}}$ with a very simple structure of which \mathcal{M} is in fact a sub-DSM ($\widehat{\mathcal{M}}$ is the completion of \mathcal{M}). As in general $\widehat{\mathcal{M}}$ preserves less relations than \mathcal{M} , the hardness of \mathcal{M} would follow. We would like the structure of $\widehat{\mathcal{M}}$ to be sufficiently simple for us to build canonically some gadgets for our hardness proof. Thus, we wish to better understand the form that elements of \mathcal{M} may take. In order to do so, we first define the canonical shop of \mathcal{M} to be the U - X shop h in \mathcal{M} , guaranteed by Proposition 16, with the property that $|h(z)|$ is maximal for each $z \in U \setminus X$. Note that this maximal h is unique, as given h_1 and h_2 of the form in Proposition 16, $h_1 \circ h_2$ is also of the required form, and further satisfies $|h_1 \circ h_2(z)| \geq |h_1(z)|, |h_2(z)|$, for all $z \in U \setminus X$.

A. Characterising reduced DSMs

Any U - X -shop in \mathcal{M} will be shown to be in the following special form, reminiscent of the form of the canonical shop.

Definition 19. We say that a shop f is in the 3-permuted form if there are a permutation ζ of $X \cap U$, a permutation χ of $X \setminus U$ and a permutation v of $U \setminus X$ such that f satisfies:

- for any y in $U \cap X$, $f(y) = \{\zeta(y)\}$;
- for any x in $X \setminus U$, $f(x) = \{\chi(x)\}$; and,
- for any u in $U \setminus X$, $f(u) = \{v(u)\} \cup X_u$, where $X_u \subseteq X \setminus U$.

Lemma 20. If a shop f satisfies $f(X) \cap (U \setminus X) = \emptyset$ then f is in the 3-permuted form.

Proof: The hypothesis forces an element of X to reach an element of X and Lemma 13 forces two elements of X to have different images. Since X is finite, there exists a permutation β of X such that for every x in X , $f(x) = \{\beta(x)\}$. Since Lemma 11 forces in particular an element of U to have at most one element of U in its image and since U is finite, it follows that there exists a permutation α of U such that for every u in U , $f(u) \cap U = \{\alpha(u)\}$ and $f^{-1}(u) \cap U = \{\alpha^{-1}(u)\}$.

It follows that there exists a permutation ζ of $U \cap X$ such that for any y in $U \cap X$, $f(y) = \{\zeta(y)\}$.

The existence of a permutation χ of $X \setminus U$ such that β is the disjoint union of χ and ζ follows. Hence, for any x in $X \setminus U$, $f(x) = \{\chi(x)\}$.

Similarly, there must also be a permutation v of $U \setminus X$ such that α is the disjoint union of v and χ . Hence, for any u in $U \setminus X$, $f(u) \cap U = \{v(u)\}$. Elements of $U \setminus X$ may however have some images in $X \setminus U$. So we get finally that for any u in $U \setminus X$, there is some $\emptyset \subseteq X_u \subseteq X \setminus U$ such that $f(u) = \{v(u)\} \cup X_u$. This proves that f is in the 3-permuted form and we are done. ■

Theorem 21 (characterisation of reduced DSM).

Let \mathcal{M} be a reduced DSM. Every shop in \mathcal{M} is in the 3-permuted form. Moreover, every U - X -shop in \mathcal{M} follows the additional requirement that the elements of $U \setminus X$ cover the set $X \setminus U$, more formally that

$$f(U \setminus X) \cap X = \bigcup_{u \in U \setminus X} X_u = X \setminus U.$$

Proof: We can now deduce easily from Lemmata 15 and 20 the form that U - X -shops in \mathcal{M} may take. It remains to prove that an arbitrary shop f in \mathcal{M} is in the 3-permuted form. Let h be the canonical shop of \mathcal{M} . It follows from Lemma 1 that $f' := h \circ f \circ h$ is a U - X -shop. Hence, f' is in the 3-permuted form. Let z in X and u in $U \setminus X$. If $f(z) \ni u$ then $f'(z) \ni u$ and f' would not be in the 3-permuted form. It follows that $f(X) \cap (U \setminus X) = \emptyset$ and appealing to Lemma 20 that f is in the 3-permuted form. ■

We do not need the following result in order to prove our main result. But surprisingly in a reduced DSM, U and X are unique. This means that we may speak of *the canonical shop of \mathcal{M}* instead its canonical U - X -shop. It also means that we can define the U - X -core of a structure \mathcal{D} without explicitly referring to U or X as the minimal substructure of \mathcal{D} which satisfy the same $\{\exists, \forall, \wedge, \vee\}$ -FO sentences.

Theorem 22. Let \mathcal{D} be a structure that is both a U - X -core and a U' - X' -core then it follows that $U = U'$ and $X = X'$.

Corollary 23 (The U - X -core with implicit U and X). Let \mathcal{D} be a finite structure. The U - X -core of \mathcal{D} is unique up to isomorphism. It is a minimal induced substructure $\tilde{\mathcal{D}}$ of \mathcal{D} , that satisfies the same $\{\exists, \forall, \wedge, \vee\}$ -FO formulae with free-variables in $\tilde{\mathcal{D}}$. Moreover, once $\tilde{\mathcal{D}}$ is fixed, there are two uniquely determined subsets U and X such that $U \cup X = |\tilde{\mathcal{D}}| \subseteq \mathcal{D}$ which are minimal within \mathcal{D} with respect to the following equivalent properties,

- \mathcal{D} has U - X -relativisation w.r.t. $\{\exists, \forall, \wedge, \vee\}$ -FO; or,
- \mathcal{D} has a U - X -shop that may act as the identity over $U \cup X$.

Remark. To simplify the presentation, we defined $\tilde{\mathcal{D}}$ as a minimal induced substructure. Considering substructures that are not necessarily induced (i.e. w.r.t. size of both the domain and the relations), we would obtain exactly the same notion (i.e. we would get minimal induced substructures). This is also the case for CSP, but it is not the case in general. For example, this is not the case for the logic $\{\exists, \forall, \wedge, =\}$ -FO, which corresponds to QCSP [9].

B. The hard DSM above \mathcal{M}

Define the completion $\widehat{\mathcal{M}}$ of \mathcal{M} to be the DSM that contains all shops in the 3-permuted form of \mathcal{M} . More precisely, the canonical shop of $\widehat{\mathcal{M}}$ is the shop \hat{h} where every set X_u is the whole set $X \setminus U$, and, for every permutation ζ of $X \cap U$, χ of $X \setminus U$ and v of $U \setminus X$, any shop in the 3-permuted form with these permutations is in $\widehat{\mathcal{M}}$. Note that by construction, \mathcal{M} is a sub-DSM of $\widehat{\mathcal{M}}$. Note also that the minimality of U and X still holds in $\widehat{\mathcal{M}}$. We will establish hardness for $\widehat{\mathcal{M}}$, whereupon hardness of \mathcal{M} follows from Theorem 2.

C. Cases II and III: NP-hardness and co-NP-hardness

We begin with Case II. We note first that $U = \{u\}$ and $|X| \geq 2$ implies $U \cap X = \emptyset$ (otherwise, we would have an $X \setminus \{u\}$ -total-shop, contradicting the minimality of X). The structure $\mathcal{K}_{|X|} \uplus \mathcal{K}_1$, the disjoint union of a clique of size $|X|$ with an isolated vertex u , has associated DSM $\widehat{\mathcal{M}}$. The complexity of $\{\exists, \wedge, \vee\}$ -FO(\mathcal{D}), according to \mathcal{D} , was fully characterised in [7] and $\{\exists, \wedge, \vee\}$ -FO($\mathcal{K}_{|X|} \uplus \mathcal{K}_1$) is NP-hard (as is any loopless graph that has at least one edge). We sketch a proof here in order to keep the paper self-contained. We observe that two structures that are homomorphically equivalent satisfy the same $\{\exists, \wedge, \vee\}$ -FO sentences. So it suffices to prove that $\{\exists, \wedge, \vee\}$ -FO($\mathcal{K}_{|X|}$) is NP-hard. If $|X| \geq 3$, then it is well known that the NP-complete problem of *graph $|X|$ -colourability* reduces to $\{\exists, \wedge\}$ -FO($\mathcal{K}_{|X|}$), otherwise $|X| = 2$ and we will need disjunction. The problem $\{\exists, \wedge, \vee\}$ -FO(\mathcal{K}_2) is NP-complete by reduction from 3-not-all-equal satisfiability (set $R_{NAE}(u, v, w) := E(u, v) \vee E(v, w)$).

For Case III, we appeal to duality. More precisely, if $\overline{\mathcal{D}}$ is the complement of \mathcal{D} , then $\text{shE}(\overline{\mathcal{D}})$ consists of exactly the inverses of the shops in $\text{shE}(\mathcal{D})$ [11] and $\{\exists, \forall, \wedge, \vee\}$ -FO($\overline{\mathcal{D}}$) is co-NP-complete iff $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) is NP-complete [13] (this

is the principle of duality, expounded ad nauseam in papers passim). It follows that we may reproduce the proof as in Case II, but with the complement structure $\overline{\mathcal{K}_{|X|} \uplus \mathcal{K}_1}$.

D. case IV: Pspace-hardness

We assume that $|U| \geq 2$ and $|X| \geq 2$ and consider two cases.

- When $U \cap X \neq \emptyset$, we show that the monoid $\widehat{\mathcal{M}}$ is a sub-DSM of a DSM \mathcal{M}' which, up to *quotienting* (to be defined shortly), is a Pspace-hard DSM with a two-element domain (dealt with in [1]).
- When $U \cap X = \emptyset$, we are unable to exhibit such a simple proof but thanks to the relative simplicity of $\widehat{\mathcal{M}}$, we provide a generic proof which is inspired from an example with a four-element domain [2].

Case 1: when $U \cap X \neq \emptyset$. Recall that if \mathcal{M} is a sub-DSM of a hard DSM $\widehat{\mathcal{M}}$ then \mathcal{M} is also hard (see Theorem 2). Moreover, in the presence of an equivalence relation on the domain of the DSM, that is a shop which maps each element to its equivalence class, we can consider the quotient (see [11]).

Lemma 24 ([11]). *Let \mathcal{D} be a structure with an equivalence relation f as a she (each element is mapped to its equivalence class). Then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) = $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}/f).*

To build $\widehat{\mathcal{M}}$ from \mathcal{M} , we added all permutations, and chose for each set $X_u = X \setminus U$. We carry on with this completion process and consider the super-DSM \mathcal{M}' which is generated by a single shop g' defined as follows:

- for every y in $X \cap U$, $g'(y) := X \cap U$; and,
- for every z in $X \Delta U$, $g'(z) := X \Delta U$, where $X \Delta U$ denotes $(X \setminus U) \cup (U \setminus X)$.

In the parlance of [11], the DSM \mathcal{M}' is a *blurred permutation subgroup* of the trivial two-element-domain DSM $\langle \begin{smallmatrix} 0 & 0 \\ 1 & 1 \end{smallmatrix} \rangle$ which is known to be Pspace-complete [1]. The result follows by Lemma 24 and Theorem 2.

Case 2: when $U \cap X = \emptyset$. We will give a structure $\widehat{\mathcal{D}}$ such that $\text{shE}(\widehat{\mathcal{D}}) = \widehat{\mathcal{M}}$. Firstly, though, given some fixed u in U and x in X , let $\mathcal{G}_{u,x}^{|U|,|X|}$ be the symmetric graph with self-loops with domain $D = U \cup X$ such that

- u and x are adjacent;
- The graph induced by X is a reflexive clique $\mathcal{K}_X^{\text{ref}}$; and,
- $U \setminus \{u\}$ and $X \setminus \{x\}$ are related via a complete bipartite graph $\mathcal{K}_{X \setminus \{x\}, U \setminus \{u\}}$.

The structure $\mathcal{G}_{u,x}^{|U|,|X|}$ and the more specific $\mathcal{G}_{1,3}^{4,5}$ are drawn in Figure 1. Denote by $E_{u,x}^{|U|,|X|}$ the binary relation of $\mathcal{G}_{u,x}^{|U|,|X|}$ and let $\widehat{\mathcal{D}}$ be the structure with a single 4-ary relation $R^{\widehat{\mathcal{D}}}$ with




Fig. 1. Main Gadget.

domain $\widehat{\mathcal{D}} = U \cup X$ specified as follows,

$$R^{\widehat{\mathcal{D}}} := \bigcup_{u \in U} \left(\left(\bigcup_{x \in X} (u, x) \times E_{u,x}^{|U|,|X|} \right) \cup \left(\bigcup_{x_1, x_2, x_3 \in X} (x_1, x_2) \times E_{u,x_3}^{|U|,|X|} \right) \right).$$

Essentially, when the first argument in a quadruple is from U , then the rest of the structure allows for the unique recovery of some $\mathcal{G}_{u,x}^{|U|,|X|}$; but if the first argument is from X then all possibilities from X for the remaining arguments are allowed. In particular, we note from the last big cup that (x_1, x_2, x_3, x_4) is a tuple of $R^{\widehat{\mathcal{D}}}$ for all quadruples x_1, x_2, x_3, x_4 in X .

Lemma 25. $\text{shE}(\widehat{\mathcal{D}}) = \widehat{\mathcal{M}}$.

Proof: Recall that, according to Theorem 21 and our assumption on U , X and $\widehat{\mathcal{M}}$, a maximal (w.r.t. sub-shop inclusion) shop f' is of the following form,

- for any x in $X \setminus U = X$, $f(x) = \{\chi(x)\}$; and,
- for any u in $U \setminus X = U$, $f(u) = \{v(u)\} \cup X$.

where χ and v are permutations of X and U , respectively.

(Backwards; $\widehat{\mathcal{M}} \subseteq \text{shE}(\widehat{\mathcal{D}})$.) It suffices to check that a maximal shop f' in $\widehat{\mathcal{M}}$ preserves $\widehat{\mathcal{D}}$. This holds by construction.

(Forwards; $\text{shE}(\widehat{\mathcal{D}}) \subseteq \widehat{\mathcal{M}}$.) We proceed by contraposition, demonstrating that $R^{\widehat{\mathcal{D}}}$ is violated by any $f \notin \widehat{\mathcal{M}}$. We consider the different ways that f might not be in $\widehat{\mathcal{M}}$.

- If f is s.t. $u \in f(x)$ for $x \in X$ and $u \in U$ then we, e.g., take $(u, x, x, x) \in R^{\widehat{\mathcal{D}}}$ but $(z, u, u, u) \notin R^{\widehat{\mathcal{D}}}$ (for any $z \in f(u)$) and we are done. It follows that $f(X) = X$.
- Assume now that f is s.t. $\{x'_1, x'_2\} \subseteq f(x)$ for $x'_1 \neq x'_2$ and $x, x'_1, x'_2 \in X$. Let $u, u' \in U$ be s.t. $u' \in f(u)$. Take $(u, x, u, x) \in R^{\widehat{\mathcal{D}}}$; $(u', x'_1, u', x'_2) \notin R^{\widehat{\mathcal{D}}}$ and we are done. It follows that f is a permutation χ on X .
- Assume now that f is s.t. $\{u'_1, u'_2\} \subseteq f(u)$ for $u'_1 \neq u'_2$ and $u, u'_1, u'_2 \in U$. Let $x, x' \in X$ be s.t. $x' \in f(x)$. Take

$(u, x, u, x) \in R^{\widehat{D}}; (u'_1, x', u'_2, x') \notin R^{\widehat{D}}$ and we are done.

It follows that f restricted to U is a permutation v on U . Hence, f is a sub-shop of a maximal shop f' from the DSM $\widehat{\mathcal{M}}$, and f belongs to $\widehat{\mathcal{M}}$ (recall that a DSM is closed under sub-shops). The result follows. ■

Proposition 26. $\{\exists, \forall, \wedge, \vee\}$ -FO(\widehat{D}) is Pspace-complete.

Proof: We begin with the observation that $\{\exists, \forall, \wedge, \vee\}$ -FO($\mathcal{G}_{u,x}^{|U|,|X|}$) is Pspace-complete (for each $u \in U$ and $x \in X$). This follows straightforwardly from the Pspace-completeness of $\{\exists, \forall, \wedge, \vee\}$ -FO($\mathcal{G}_{1,3}^{2,2}$), the simplest gadget which is depicted on Figure 1a. These gadgets $\mathcal{G}_{u,x}^{|U|,|X|}$ agree on all equality-free sentences – even ones involving negation – as there is a strong homomorphism from $\mathcal{G}_{u,x}^{|U|,|X|}$ to $\mathcal{G}_{1,3}^{2,2}$; see the Homomorphism Theorem in [14].

We will prove that $\{\exists, \forall, \wedge, \vee\}$ -FO($\mathcal{G}_{1,3}^{2,2}$) is Pspace-hard, by reduction from QCSP(\mathcal{B}_{NAE}) – *quantified 3-not-all-equal satisfiability* (see [15]). Recall that we may assume w.l.o.g. that \forall variables are relativised to U and that \exists -variables are relativised to X , by Theorem 8. Let φ be an instance of QCSP(\mathcal{B}_{NAE}). We reduce φ to a (relativised) instance ψ of $\{\exists, \forall, \wedge, \vee\}$ -FO($\mathcal{G}_{1,3}^{2,2}$). The reduction goes as follows:

- an existential variable $\exists x$ of φ is replaced by an existential variable $\exists v_x \in X$ in ψ ;
- a universal variable $\forall u$ of φ is replaced by $\forall u \in U, \exists v_u \in X, E(u, v_u)$ in ψ ; and,
- every clause $C_i := R(\alpha, \beta, \gamma)$ in φ is replaced by $\forall c_i \in U, E(c_i, v_\alpha) \vee E(c_i, v_\beta) \vee E(c_i, v_\gamma)$ in ψ .

The truth assignment is read from \exists choices in X for the variables v : we arbitrarily see one value in X as true and the other as false. It is not relevant which one is which for the problem not-all-equal satisfiability, we only need to ensure that no three variables involved in a clause can get the same value. The $\forall c_i \in U$ acts as a conjunction, enforcing “one of $v_\alpha, v_\beta, v_\gamma$ is true” and “one of $v_\alpha, v_\beta, v_\gamma$ is false”. This means that at least one in three has a different value.

Now, we can prove that $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}') is Pspace-complete by substituting $R(u_0, x_0, u, v)$ for each instance of $E(u, v)$ in the previous proof, and by quantifying the sentence so-produced with the prefix $\forall u_0 \in U, \exists x_0 \in X$, once u_0 and x_0 are chosen, play proceeds as above but in the copy $\mathcal{G}_{u_0, x_0}^{|U|,|X|}$, and the result follows. ■

V. THE COMPLEXITY OF THE META-PROBLEM

The $\{\exists, \forall, \wedge, \vee\}$ -FO(σ) meta-problem takes as input a finite σ -structure \mathcal{D} and answers L, NP-complete, co-NP-complete or Pspace-complete, according to the complexity of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}). The principle result of this section is that this problem is NP-hard even for some fixed and finite signature σ_0 .²

Note that one may determine if a given shop f is a she of a structure \mathcal{D} in, say, quadratic time in $|D|$. Since we are

²For now, σ_0 consists of two binary predicates and three monadic predicates. The monadic predicates are for convenience, but it is not clear whether a single binary predicate suffices.

not interested here in distinguishing levels within P, we will henceforth consider such a test to be a basic operation. We begin with the most straightforward case.

Proposition 27. On input \mathcal{D} , the question “is $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) in L?” is in P.

Proof: By Theorem 3, we need to check whether \mathcal{D} has both an A-she and an E-she. In this special case, it suffices to test for each u, x in D , if the following $\{u\}$ - $\{x\}$ -she f preserves \mathcal{D} (this shop was denoted by $\forall_u \exists_x$ in [1], [8]): $f(u) := D$ and $f^{-1}(x) := D$. ■

Proposition 28. For some fixed and finite signature σ_0 , on input of a σ -structure \mathcal{D} , the question “is $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) in NP (respectively, NP-complete, in co-NP, co-NP-complete)?” is NP-complete.

Proof: The four variants are each in NP. For the first, one guesses and verifies that \mathcal{D} has an A-she, for the second, one further checks that there is no $\forall_u \exists_e$ -she (see the proof of Proposition 27). Similarly for the third, one guesses and verifies that \mathcal{D} has an E-she; and, for the fourth, one further checks that there is no $\forall_u \exists_e$ -she. The result then follows from Theorem 3.

For NP-hardness we will address the first problem only. The same proof will work for the second (for the third and fourth, recall that a structure \mathcal{D} has an A-she iff its complement $\overline{\mathcal{D}}$ has an E-she). We reduce from *graph 3-colourability*. Let \mathcal{G} be an undirected graph with vertices $V := \{v_1, v_2, \dots, v_s\}$. We will build a structure $\mathcal{S}_{\mathcal{G}}$ over the domain D which consists of the disjoint union of “three colours” $\{0, 1, 2\}$, u , and the “vertices” from V .

The key observation is that there is a structure \mathcal{G}_V whose class of shes $\text{shE}(\mathcal{G}_V)$ is generated by the following A-shop:

$$f_V := \begin{array}{c|c} 0 & 0 \\ \hline 1 & 1 \\ \hline 2 & 2 \\ \hline u & 0, 1, 2, u, v_1, \dots, v_s \\ \hline v_1 & 0, 1, 2 \\ \hline v_2 & 0, 1, 2 \\ \hline \vdots & \vdots \\ \hline v_s & 0, 1, 2 \end{array}$$

The existence of such a \mathcal{G}_V is in fact guaranteed by the Galois connection, fully given in [11], but that may require relations of unbounded arity, and we wish to establish our result for a fixed signature. So we will appeal to Lemma 29, below, for a σ_V -structure \mathcal{G}_V with the desired class of shes, where the signature σ_V consists of one binary relation and three monadic predicates. The signature σ_0 will be σ_V together with a binary relational symbol E .

The structure $\mathcal{S}_{\mathcal{G}}$ is defined as in \mathcal{G}_V for symbols in σ_V , and for the additional binary symbol E , as the edge relation of the instance \mathcal{G} of 3-colourability together with a clique \mathcal{K}_3 for the colours $\{1, 2, 3\}$. By construction, the following holds.

- Any she g of $\mathcal{S}_{\mathcal{G}}$ will be a subshop of f_V .
- Restricting such a shop g to V provides a set of mutually consistent 3-colourings: i.e. we may pick arbitrarily a

colour from $g(v_i)$ to get a 3-colouring \tilde{g} .

- Conversely, a 3-colouring \tilde{g} induces a sub-shop g of f_V : set g as f_V over elements from $\{0, 1, 2, u\}$ and as \tilde{g} over V .

This proves that graph 3-colourability reduces to the meta-question “is $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) in NP”. ■

Note that it follows from the given proof that the meta-problem itself is NP-hard. To see this, we take the structure $\mathcal{S}_{\mathcal{G}}$ from the proof of Proposition 28 and ask which of the four classes L, NP-complete, co-NP-complete or Pspace-complete the corresponding problem belongs to. If the answer is NP-complete then \mathcal{G} was 3-colourable; otherwise the answer is Pspace-complete and \mathcal{G} was not 3-colourable.

Lemma 29. *Let σ_V be a signature involving one binary relations E' and three monadic predicates One, Two and Three. There is a σ_V -structure \mathcal{G}_V such that $\text{shE}(\mathcal{G}_V) = \langle f_V \rangle$.*

VI. CONCLUSION

We complete the complexity classification of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) started in [1], [2], [8] and prove our conjecture from [2], [8].

In order to prove this result, we introduce the notion of U - X -core, a generalisation of the well-known concept of core that is so useful in the context of CSP, and which allows us to dispense with working out complexities from within the lattice of DSMs, and escape combinatorial explosion. The notion of a U - X -core is characterised logically both in terms of relativisation properties of sentences of $\{\exists, \forall, \wedge, \vee\}$ -FO; and, algebraically in terms of preservation by a U - X -shop. This means that each of the four cases can be formulated as in Table I.

TABLE I
REFORMULATIONS OF THE TETRACHOTOMY

Case	I	
Complexity	L	
A-shop	yes	
E-shop	yes	
U - X -core	$ U = 1, X = 1$	
Relativises into	$\{\wedge, \vee\}$ -FO	
Dual to case	self-dual	
Case	II	III
Complexity	NP-complete	co-NP-complete
A-shop	yes	no
E-shop	no	yes
U - X -core	$ U = 1, X \geq 2$	$ U \geq 2, X = 1$
Relativises into	$\{\exists, \wedge, \vee\}$ -FO	$\{\forall, \vee, \wedge\}$ -FO
Dual to case	III	II
Case	IV	
Complexity	Pspace-complete	
A-shop	no	
E-shop	no	
U - X -core	$ U \geq 2, X \geq 2$	
Relativises into	$\{\exists, \forall, \wedge, \vee\}$ -FO	
Dual to case	self-dual	

We believe that U - X -core could be of interest for other purposes. Firstly, since equality may be propagated out for QCSP, computing the U - X -core could be an effective pre-processing step for a QCSP-solver, which could reduce the

branching factor during search.³ Secondly, though it is not the right notion of a core for QCSP, it could help with the classification, for example for graphs and digraphs, for which only very partial classifications are known [16], [17]. The key property here being not only that we could shrink the domain, but that we could relativise quantifiers, in particular the universal ones.⁴

The algebraic machinery that we have set up in the course of our investigation is particularly well-behaved and pleasant to use and could be of further use. It will be possible to extend it slightly to characterise, and show the decidability of, problem containment for $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{D}) (this was investigated for QCSP in [18]). Finally, we only considered the usual decision problem associated with $\{\exists, \forall, \wedge, \vee\}$ -FO, and it could be that a number of other interesting combinatorial questions associated with $\{\exists, \forall, \wedge, \vee\}$ -FO could be solved using this nice algebraic setting.

REFERENCES

- [1] F. Madelaine and B. Martin, “The complexity of positive first-order logic without equality,” in *LICS*. IEEE Computer Society, 2009.
- [2] B. Martin and J. Martin, “The complexity of positive first-order logic without equality II: The four-element case,” in *CSL 2010*, ser. Incs, A. Dawar and H. Veith, Eds., vol. 6247. Springer, 2010, pp. 426–438.
- [3] T. Schaefer, “The complexity of satisfiability problems,” in *STOC*, 1978.
- [4] A. A. Bulatov, “A dichotomy theorem for constraint satisfaction problems on a 3-element set,” *J. ACM*, vol. 53, no. 1, pp. 66–120, 2006.
- [5] H. Chen, “The complexity of quantified constraint satisfaction: Collapsibility, sink algebras, and the three-element case,” *SIAM J. Comput.*, vol. 37, no. 5, pp. 1674–1701, 2008.
- [6] F. Börner, A. A. Bulatov, H. Chen, P. Jeavons, and A. A. Krokhin, “The complexity of constraint satisfaction games and QCSP,” *Inf. Comput.*, vol. 207, no. 9, pp. 923–944, 2009.
- [7] B. Martin, “First order model checking problems parameterized by the model,” in *CiE 2008, LNCS 5028*, 2008, pp. 417–427.
- [8] F. R. Madelaine and B. Martin, “The complexity of positive first-order logic without equality,” *CoRR*, vol. abs/1003.0802, 2010, to appear in ACM TOCL.
- [9] F. Madelaine and B. Martin, “The preservation properties of positive Horn logic,” 2009, manuscript.
- [10] F. Börner, “Total multifunctions and relations,” in *AAA60: Workshop on General Algebra, Dresden, Germany*, 2000.
- [11] B. Martin, “The lattice structure of sets of surjective hyper-operations,” in *CP*, 2010, pp. 368–382.
- [12] N. Lynch, “Log space recognition and translation of parenthesis languages,” *J. ACM*, vol. 24, pp. 583–590, 1977.
- [13] B. Martin, “Model checking positive equality-free FO: Boolean structures and digraphs of size three,” *CoRR*, vol. abs/0808.0647, 2008.
- [14] H. B. Enderton, *A Mathematical Introduction to Logic*. AP, 1972.
- [15] C. Papadimitriou, *Computational Complexity*. Addison-Wesley, 1994.
- [16] B. Martin and F. R. Madelaine, “Towards a trichotomy for quantified H-coloring,” in *CiE 2006, LNCS 3988*, 2006, pp. 342–352.
- [17] B. Martin and F. Madelaine, “Quantified constraints on directed graphs,” in *London Algorithmics and Stringology 2006*, ser. Texts in Algorithmics, M. M. J. Daykin and K. Steinhilf, Eds. King’s College publication, 2007, ISBN: 9781904987413.
- [18] H. Chen, F. R. Madelaine, and B. Martin, “Quantified constraints and containment problems,” in *LICS*. IEEE Computer Society, 2008.

³One of the key features of CSP-solvers, filtering, though polynomial, is typically expensive for QCSP because of the branching factor. Computing the U - X -core is exponential but as it would be done only once, it could end up being effective.

⁴It is not difficult to restrict \exists ’s play in order to fit the design of a reduction, but it is hard to do so for the universal player, and is one important technical hurdle. In fact this is a challenge even in the case of $\{\exists, \forall, \wedge, \vee\}$ -FO: contrast the simplicity of our hardness proof from Section IV which use relativisation, with the relatively involved proof used in the four-element case [2].