

THE ADJACENT FRAGMENT AND QUINE'S LIMITS OF DECISION

BARTOSZ BEDNARCZYK, DAUMANTAS KOJELIS, AND IAN PRATT-HARTMANN

Abstract. We introduce the *adjacent fragment* \mathcal{AF} of first-order logic, obtained by restricting the sequences of variables occurring as arguments in atomic formulas. The adjacent fragment generalizes (after a routine renaming) the two-variable fragment of first-order logic as well as the so-called fluted fragment. We show that the adjacent fragment has the finite model property, and that the satisfiability problem for its k -variable sub-fragment is in $(k-1)$ -NEXPTIME. Using known results on the fluted fragment, it follows that the satisfiability problem for the whole adjacent fragment is TOWER-complete. We additionally consider the effect of the adjacency requirement on the well-known guarded fragment of first-order logic, whose satisfiability problem is 2EXPTIME-complete. We show that the satisfiability problem for the intersection of the adjacent and guarded adjacent fragments remains 2EXPTIME-hard. Finally, we show that any relaxation of the adjacency condition on the allowed order of variables in argument sequences yields a logic whose satisfiability and finite satisfiability problems are undecidable.

§1. Introduction. The quest to find fragments of first-order logic for which satisfiability is algorithmically decidable has been a central undertaking of mathematical logic since the appearance of D. Hilbert and W. Ackermann's *Grundzüge der theoretischen Logik* [19, 20] almost a century ago. The best-known of these fragments belong to just three families: (i) *quantifier prefix* fragments [11], where we are restricted to formulas in prenex form with a specified quantifier sequence; (ii) *two-variable* logics [17], where the only logical variables occurring as arguments of predicates are x_1 and x_2 ; and (iii) *guarded* logics, where either quantifiers or negated formulas are relativized by atomic formulas featuring all the free variables in their scope [1, 4]. There is, however, a fourth family of first-order fragments for which satisfiability is decidable, but which has languished in relative obscurity. The fragments of this family are defined by restricting the allowed sequences of variables occurring as arguments in atomic formulas, an idea dating back W. Quine's *homogeneous m -adic formulas* [34]. Such *argument-sequence fragments*, as we might call them, include the *fluted fragment* [33], the *ordered fragment* [18] and the *forward fragment* [5]. In this paper, we identify a new argument-sequence fragment, the *adjacent fragment*, which includes the fluted, ordered and forward fragments, and subsumes, in a sense we make precise, the two-variable fragment. We show that the satisfiability problem for the adjacent fragment is decidable, and obtain bounds on its complexity. Finally, we show that the adjacent fragment is maximal among the argument-sequence fragments whose satisfiability and finite satisfiability problems are decidable.

Bartosz Bednarczyk was supported by the ERC Consolidator Grant No. 771779 (DeciGUT). Ian Pratt-Hartmann was supported by the NCN grant 2018/31/B/ST6/03662.

To explain how restrictions on argument sequences work, we consider presentations of first-order logic over purely relational signatures, employing individual variables from the alphabet $\{x_1, x_2, x_3, \dots\}$. Any atomic formula in this logic has the form $p(\bar{x})$, where p is a predicate of arity m (possibly 0), and \bar{x} a word of length m over the alphabet of variables. Call a first-order formula φ *index-normal* if any occurrence of a quantifier binding a variable x_k has as its scope a Boolean combination of formulas that either (i) are atomic with free variables among x_1, \dots, x_k , or (ii) have as their major connective a quantifier binding x_{k+1} . By re-indexing variables, any first-order formula can easily be written as a logically equivalent index-normal formula. In the *fluted fragment*, as defined by W. Purdy [33, Sec. 3], we confine attention to index-normal formulas, but additionally insist that any atom occurring in a context in which x_k is available for quantification have the form $p(x_{k-m+1} \cdots x_k)$, i.e. $p(\bar{x})$ with \bar{x} being a *suffix* of $x_1 \cdots x_k$. In the *ordered fragment*, due to A. Herzig [18, Sec. 2], by contrast, we insist that \bar{x} be a *prefix* of $x_1 \cdots x_k$. In the *forward fragment*, due to B. Bednarczyk [5, Sec. 3.1], we insist only that \bar{x} be an *infix* (i.e. a factor) of $x_1 \cdots x_k$. All these logics have the finite model property, and hence are decidable for satisfiability.

We denote the fluted fragment by \mathcal{FL} , and the sub-fragment of \mathcal{FL} involving at most k variables (free or bound) by \mathcal{FL}^k . It is known that the satisfiability problem for \mathcal{FL}^k is in $(k-2)$ -NEXPTIME for all $k \geq 3$, and $\lfloor k/2 \rfloor$ -NEXPTIME-hard for all $k \geq 2$ [31, Thm. 3.2 & Thm. 4.2]. Thus, satisfiability for the whole of \mathcal{FL} is TOWER-complete, in the system of trans-elementary complexity classes due to S. Schmitz [36, Sec. 3.1.1]. By contrast, the satisfiability problem for the ordered fragment is PSPACE-complete [18] (see also a related result by R. Jaakkola [21, Thm. 13]). On the other hand, the apparent liberalization afforded by the forward fragment yields no useful increase in expressive power, and there is a polynomial-time, satisfiability-preserving reduction of the forward fragment to the fluted fragment [7, p. 182]. The term “fluted” originates with Quine [35], and presumably invites us to imagine the atoms in formulas aligned in such a way that the variables form columns. (However, it is unclear that Quine had in mind the fragment now generally referred to as the fluted fragment; for a brief historical discussion see, [29, p. 221].) Note that none of these fragments can state that a relation is reflexive or symmetric, as can be easily established using a game-theoretic argument [7, Sec. 3].

Say that a word \bar{x} over the alphabet $\{x_1, \dots, x_k\}$ is *adjacent* if the indices of neighbouring letters differ by at most 1. For example, $x_3x_2x_1x_2x_2x_2x_3x_4x_3$ is adjacent, but $x_1x_3x_2$ is not. The *adjacent fragment* is analogous to the fluted, ordered and forward fragments, but we allow any atom $p(\bar{x})$ to occur in a context where x_k is available for quantification as long as \bar{x} is an adjacent word over $\{x_1, \dots, x_k\}$ (see Sec. 2 for a formal definition). As a simple example, the formula

$$(1) \quad \forall x_1 \forall x_2 \forall x_3 \exists x_4 \forall x_5 (p(x_1x_2x_3x_2x_3x_4x_5) \rightarrow p(x_1x_2x_3x_4x_3x_4x_5)).$$

is in the adjacent fragment. (In fact, it is a validity, as can be seen by assigning x_4 the same value as x_2 .) We denote the adjacent fragment by \mathcal{AF} , and the sub-fragment of \mathcal{AF} involving at most k variables (free or bound) by \mathcal{AF}^k . Evidently, \mathcal{AF} includes the fluted, ordered and forward fragments; the inclusion is strict,

since the formulas $\forall x_1 r(x_1x_1)$ and $\forall x_1x_2(r(x_1x_2) \rightarrow r(x_2x_1))$, stating that r is reflexive and symmetric, respectively, are in \mathcal{AF} . As we show in the sequel (Theorem 4), any formula of the two-variable fragment may be translated to a logically equivalent formula of \mathcal{AF} . Hence, a number of other well-known logics can be translated naturally into the adjacent fragment, including the system of basic multimodal propositional logic K (under the standard translation), a great many description logics [3], and even polyadic extensions of multimodal logic [14, Sec. 1.5].

Our principal result is that \mathcal{AF} has the finite model property, and that the satisfiability problem for \mathcal{AF}^k is in $(k-1)$ -NEXPTIME for all $k \geq 2$. The proof follows the same basic strategy as employed for \mathcal{FL}^k in [31]: the (finite) satisfiability problem for \mathcal{AF}^{k+1} is reduced, with exponential blow-up, to that for \mathcal{AF}^k . The result then follows from the fact that \mathcal{AF}^2 is subsumed by the two-variable fragment, FO^2 , which has the finite model property, and for which satisfiability is known to be in NEXPTIME [16]. On the other hand, \mathcal{AF}^k includes \mathcal{FL}^k , whence the above-mentioned lower bounds for the latter carry over: the satisfiability problem for \mathcal{AF}^k is $\lfloor k/2 \rfloor$ -NEXPTIME-hard for all $k \geq 2$ (and NPTIME-hard for $k \in \{0, 1\}$). We remark that, using techniques similar to those employed in [31], we can in fact shave one exponential off the upper bounds for \mathcal{AF}^k ($k \geq 3$) when the equality predicate is disallowed; in the interests of simplicity, we leave this as an exercise to the interested reader. We additionally consider the *guarded adjacent fragment* \mathcal{GA} , defined as the intersection of the adjacent fragment, \mathcal{AF} , and the guarded fragment, \mathcal{GF} , due to H. Andréka, J. van Benthem, and I. Németi [1, Sec. 4.1]. The satisfiability problem for \mathcal{GF} is 2EXPTIME-complete, as shown by E. Grädel [15, Thm. 4.4]. We show in the sequel that the satisfiability problem for \mathcal{GA} remains 2EXPTIME-hard. We finish with a pair of results on the expressiveness of the adjacent fragment. First, we show that this fragment subsumes, in a sense we make precise, the two-variable fragment, with a converse subsumption holding for signatures featuring predicates of arity at most two. Second, we consider liberalizations of the adjacent fragment, in which the palette of permitted variable sequences is further extended. We show that the fragment \mathcal{AF} is a *maximal* argument-sequence fragment for which satisfiability (or finite satisfiability) is decidable.

The structure of the paper is as follows. Sec. 2 defines the fragments considered in this paper and establishes the notation used throughout. Sec. 3 is devoted to the combinatorics of words, and presents a pair of results (Lemmas 2 and 3), which form the basis of the following two sections. Sec. 4 establishes upper bounds on the complexity of satisfiability for the sub-fragment of \mathcal{AF}^k *without equality*; concentrating on the equality-free case simplifies the combinatorics, thus bringing the key technical ideas into sharper focus. Sec. 5 then extends these results to the logic with equality. Sec. 6 gives the advertised lower complexity bound for the guarded adjacent fragment. Sec. 7 establishes the observations on expressive power mentioned in the previous paragraph. The results of Secs. 4 and 6, concerning the adjacent fragment without equality, were first announced in the conference paper [8]. The present article provides full proofs, and extends our results to the full adjacent fragment with equality.

§2. Preliminaries. Let m and k be non-negative integers. For any integers i and j , we write $[i, j]$ to denote the set of integers h such that $i \leq h \leq j$. A function $f: [1, m] \rightarrow [1, k]$ is *adjacent* if $|f(i+1) - f(i)| \leq 1$ for all $i \in [1, m-1]$. We write \mathbf{A}_k^m to denote the set of adjacent functions $f: [1, m] \rightarrow [1, k]$. Since $[1, 0] = \emptyset$, we have $\mathbf{A}_k^0 = \{\emptyset\}$, and $\mathbf{A}_0^m = \emptyset$ if $m > 0$. Let A be a non-empty set. Regarding A as an alphabet, a word \bar{a} over the alphabet A is simply a tuple of elements from A ; we alternate freely in the sequel between these two ways of speaking, as the context requires. Accordingly, we take A^k to denote the set of words over A having length exactly k , and A^* , the set of all finite words over A . If $\bar{a} = a_1 \cdots a_k$, we write $|\bar{a}| = k$ for the length of \bar{a} , and $\tilde{a} = a_k \cdots a_1$ for the reversal of \bar{a} . Any function $f: [1, m] \rightarrow [1, k]$ (adjacent or not) induces a natural map from A^k to A^m defined by $\bar{a}^f = a_{f(1)} \cdots a_{f(m)}$. If $f \in \mathbf{A}_k^m$ (i.e. if f is adjacent), we may think of \bar{a}^f as the result of a ‘going for a stroll’ on the tuple \bar{a} , starting at the element $a_{f(1)}$, and moving left, right, or remaining stationary according to the sequence of values $f(i+1) - f(i)$ (for $1 \leq i < m$).

For any $k \geq 0$, denote by \mathbf{x}_k the fixed word $x_1 \cdots x_k$ (if $k = 0$, this is the empty word). A k -atom is an expression $p(\mathbf{x}_k^f)$, where p is a predicate of some arity m , and $f: [1, m] \rightarrow [1, k]$. Thus, in a k -atom, each argument is a variable chosen from \mathbf{x}_k . If f is adjacent, we speak of an *adjacent k -atom*. Thus, in an adjacent k -atom, the indices of neighbouring arguments differ by at most one. The equality predicate is allowed when $m = 2$. Proposition letters (predicates of arity $m = 0$) count as (adjacent) k -atoms for all $k \geq 0$, taking f to be the empty function. When $k = 0$, we perforce have $m = 0$, since otherwise, there are no functions from $[1, m]$ to $[1, k]$; thus the 0-atoms are precisely the proposition letters. When $k \leq 2$, the adjacency requirement is vacuous, and we prefer to speak simply of k -atoms.

We define the sets of first-order formulas $\mathcal{AF}^{[k]}$ by simultaneous structural induction for all $k \geq 0$:

1. every adjacent k -atom is in $\mathcal{AF}^{[k]}$;
2. $\mathcal{AF}^{[k]}$ is closed under Boolean combinations;
3. if φ is in $\mathcal{AF}^{[k+1]}$, then $\exists x_{k+1} \varphi$ and $\forall x_{k+1} \varphi$ are in $\mathcal{AF}^{[\ell]}$ for all $\ell \geq k$.

Now let $\mathcal{AF} = \bigcup_{k \geq 0} \mathcal{AF}^{[k]}$ and define \mathcal{AF}^k to be the set of formulas of \mathcal{AF} featuring no variables other than x_1, \dots, x_k , free or bound. We call \mathcal{AF} the *adjacent fragment* and \mathcal{AF}^k the *k -variable adjacent fragment*. Note that formulas of \mathcal{AF} contain no individual constants or function symbols; however, they may contain equality. The primary objects of interest here are the languages \mathcal{AF} and \mathcal{AF}^k ; the sets of formulas $\mathcal{AF}^{[k]}$ will make only occasional appearances in the sequel. Thus, for example, the formula (1) is in \mathcal{AF}^k if and only if $k \geq 5$, but it is in $\mathcal{AF}^{[k]}$ for all $k \geq 0$. On the other hand, the *quantifier-free* formulas of $\mathcal{AF}^{[k]}$ and \mathcal{AF}^k are the same. A simple structural induction establishes that $\mathcal{AF}^{[k]} \subseteq \mathcal{AF}^{[\ell]}$ for all $k \leq \ell$.

We silently assume the variables $\mathbf{x}_k := x_1 \cdots x_k$ to be ordered in the standard way. That is: if φ is a formula of $\mathcal{AF}^{[k]}$, \mathfrak{A} a structure interpreting its signature, and $\bar{a} := a_1 \cdots a_k \in A^k$, we say simply that \bar{a} *satisfies* φ in \mathfrak{A} , and write $\mathfrak{A} \models \varphi[\bar{a}]$ to mean that \bar{a} satisfies φ in \mathfrak{A} under the assignment $x_i \leftarrow a_i$ (for all $1 \leq i \leq k$). (This does not necessarily mean that each of the variables of \mathbf{x}_k actually appears

in φ .) If φ is true under all assignments in all structures, we write $\models \varphi$; the notation $\varphi \models \psi$ means the same as $\models \varphi \rightarrow \psi$ (i.e. variables are consistently instantiated in φ and ψ). The notation $\varphi(\bar{v})$, where $\bar{v} := v_1 \cdots v_k$ are variables (chosen from among x_1, x_2, \dots), will always be used to denote the formula that results from substituting v_i for x_i ($1 \leq i \leq k$) in φ , rather than to indicate the order in which elements of some structure are to be assigned to variables. If φ is any formula, $\text{fv}(\varphi)$ denotes the set of free variables of φ . A *sentence* is a formula with no free variables. Necessarily, all formulas of $\mathcal{AF}^{[0]}$ are sentences. For a sentence φ we write simply $\mathfrak{A} \models \varphi$ to mean that φ is true in \mathfrak{A} . We call the set of predicates used in φ *the signature of φ* , denoted $\text{sig}(\varphi)$.

We adapt the standard notion of (atomic) k -types for the fragments studied here. Fix some non-logical relational signature τ (i.e. not containing the equality predicate). An *adjacent k -literal* over τ is an *adjacent k -atom* or its negation, featuring a predicate in $\tau \cup \{=\}$. An *adjacent k -type* over τ is a maximal consistent set of adjacent k -literals over τ . Reference to τ is suppressed where clear from context. We use the letters ζ, η and ξ to range over adjacent k -types for various k . We denote by Atp_k^τ the set of all adjacent k -types over τ . For finite τ , we identify members of Atp_k^τ with their conjunctions, and treat them as (quantifier-free) \mathcal{AF}^k -formulas, writing ζ instead of $\bigwedge \zeta$. Given a pair of integers i, j ($1 \leq i \leq j \leq k$), we write $\zeta|_{[i,j]}$ for the set η obtained by deleting literals in ζ that feature variables outside the range $[i, j]$. It is evident that (after a shift in indices) η is a $(j-i+1)$ -type. We write ζ^+ for the quantifier-free $\mathcal{AF}^{\ell+1}$ -formula $\zeta(x_2, \dots, x_{k+1})$ obtained by incrementing the index of each variable. When $k \leq 2$, the adjacency requirement is vacuous, and in this case we shall simply speak of *k -types*. Every quantifier-free \mathcal{AF}^k -formula χ is thus logically equivalent to a disjunction of adjacent k -types, as may be seen by writing χ in disjunctive normal form. In particular, if χ is satisfiable, then there is an adjacent k -type which entails it. If \mathfrak{A} is a τ -structure and \bar{a} a k -tuple of elements from A , there is a unique adjacent k -type ζ such that $\mathfrak{A} \models \zeta[\bar{a}]$; we denote this adjacent k -type by $\text{atp}^\mathfrak{A}[\bar{a}]$, and call it the *adjacent type of \bar{a} in \mathfrak{A}* . It is not required that the elements of \bar{a} be distinct; note however that any (in)equality literals occurring in adjacent types must themselves be adjacent. For instance, $x_5 = x_6$ or $x_5 \neq x_6$ may occur in an adjacent type, but not $x_4 = x_6$ or $x_4 \neq x_6$.

The following derivative notions relating to adjacent types will feature in the sequel. Call an adjacent k -literal *covering* if it features all of the variables in \mathbf{x}_k , i.e. if it has the form $p(\mathbf{x}_k^f)$ or $\neg p(\mathbf{x}_k^f)$ with p of arity m and $f \in \mathbf{A}_k^m$ surjective. Define an *incremental k -type* over τ to be a maximal consistent set ι of covering adjacent k -literals over τ . If ξ is an adjacent k -type, then the *increment of ξ* , denoted $\partial\xi$, is the (unique) incremental k -type included in ξ . If \mathfrak{A} interprets τ and \bar{c} is a k -tuple over \mathfrak{A} , then the *incremental type of \bar{c} in \mathfrak{A}* is the (unique) incremental type ι such that $\mathfrak{A} \models \iota[\bar{c}]$; we write $\text{itp}^\mathfrak{A}[\bar{c}]$ to denote ι . Suppose now that $\bar{c} := a\bar{a}b$, $\text{atp}^\mathfrak{A}[a\bar{a}] = \zeta$, $\text{atp}^\mathfrak{A}[\bar{a}b] = \eta$ and $\text{itp}^\mathfrak{A}[\bar{c}] = \iota$. It should be obvious that, writing ξ for $\text{atp}^\mathfrak{A}[\bar{c}]$, we have $\xi = \zeta \cup \eta^+ \cup \iota$, and, moreover, $\iota = \partial\xi$.

The following lemma establishes a normal form for $\mathcal{AF}^{\ell+1}$ -sentences, which simplifies the decision procedures discussed in Secs. 4 and 5.

LEMMA 1. *Let φ be a sentence of $\mathcal{AF}^{\ell+1}$, where $\ell \geq 1$. We can compute, in polynomial time, an $\mathcal{AF}^{\ell+1}$ -formula ψ satisfiable over the same domains as φ , of the form*

$$(2) \quad \psi := \bigwedge_{i \in I} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_i \wedge \forall \mathbf{x}_{\ell+1} \beta,$$

where I is a finite index set, and the formulas γ_i and β are quantifier-free; moreover, if φ is equality-free, then so is ψ .

PROOF. If the sentence φ is quantifier-free, then it is a formula of the propositional calculus, and the result is easily obtained by adding vacuous quantification. Otherwise, write $\varphi_0 = \varphi$, and let $\theta := Qx_{k+1} \chi$ be a subformula of φ , where $Q \in \{\forall, \exists\}$, such that χ is quantifier-free. Writing $\bar{\exists} = \forall$ and $\bar{\forall} = \exists$, let p be a new predicate of arity k , let φ_1 be the result of replacing θ in φ_0 by the atom $p(\mathbf{x}_k)$, and let ψ_1 be the formula

$$\forall \mathbf{x}_k Qx_{k+1} (p(\mathbf{x}_k) \rightarrow \chi) \wedge \forall \mathbf{x}_k \bar{Q}x_{k+1} (\chi \rightarrow p(\mathbf{x}_k)).$$

It is immediate that $\varphi_1 \wedge \psi_1 \models \varphi_0$. Conversely, if $\mathfrak{A} \models \varphi_0$, then we may expand \mathfrak{A} to a model \mathfrak{A}' of $\varphi_1 \wedge \psi_1$ by taking $p^{\mathfrak{A}'}$ to be the set of k -tuples \bar{a} such that $\mathfrak{A} \models \theta[\bar{a}]$. Evidently, φ_1 is a sentence of $\mathcal{AF}^{\ell+1}$. Processing φ_1 in the same way, and proceeding similarly, we obtain a set of formulas $\varphi_2, \dots, \varphi_m$ and ψ_2, \dots, ψ_m , with φ_m quantifier-free and φ_0 satisfiable over the same domains as $\psi_1 \wedge \dots \wedge \psi_m \wedge \varphi_m$. Since φ_m is a sentence, it is a formula of the propositional calculus. By moving φ_m inside one of the quantified formulas, re-indexing variables and re-ordering conjuncts, we obtain a formula ψ of the form (2). \dashv

We refer to any \mathcal{AF}^k -sentence having the form (2) as a *normal-form* formula.

The following notation will be useful. If χ is any quantifier-free $\mathcal{AF}^{\ell+1}$ -formula, we denote by χ^{-1} the formula $\chi(x_{\ell+1}, \dots, x_1)$ obtained by simultaneously replacing each variable x_h by $x_{\ell-h+2}$ ($1 \leq h \leq \ell+1$); and we denote by $\hat{\chi}$ the formula $\chi \wedge \chi^{-1}$. Obviously χ^{-1} and $\hat{\chi}$ are also in $\mathcal{AF}^{\ell+1}$.

§3. Primitive generators of words. The upper complexity bounds obtained below depend on an observation concerning the combinatorics of words, which may be of independent interest. For words $\bar{a}, \bar{c} \in A^*$ with $|\bar{a}| = k$ and $|\bar{c}| = m$, say that \bar{a} *generates* \bar{c} if $\bar{c} = \bar{a}^f$ for some *surjective* function $f \in \mathbf{A}_k^m$. As explained above, it helps to think of \bar{a}^f as the sequence of letters encountered on an m -step ‘stroll’ backwards and forwards on the tuple \bar{a} , with $f(i)$ giving the index of our position in \bar{a} at the i th step. The condition that f is adjacent ensures that we never change position by more than one letter at a time; the condition that f is surjective ensures that we visit every position of \bar{a} . We may depict f as a piecewise linear graph, with the generated word superimposed on the abscissa and the generating word on the ordinate (Fig. 1). We refer to any maximal interval $[i, j] \subseteq [1, m]$ over which $f(h+1) - f(h)$ is constant (for $i \leq h < j$) as a *leg* of f . Thus, the legs correspond to the straight-line segments in the graph of f . A leg is *increasing*, *flat* or *decreasing* according as $f(h+1) - f(h)$ is 1, 0 or -1.

Every word \bar{a} generates both itself and its reversal, \tilde{a} . Moreover, if \bar{a} generates \bar{c} , then $|\bar{c}| \geq |\bar{a}|$, by the surjectivity requirement. In fact, \bar{a} and \tilde{a} are the only




FIGURE 1. Generation of abcbaaaadefedadedfbabf from cbadefba.

words of length $|\bar{a}|$ generated by \bar{a} . Finally, generation is transitive: if \bar{a} generates \bar{b} and \bar{b} generates \bar{c} , then \bar{a} generates \bar{c} . We call \bar{a} *primitive* if it is not generated by any word shorter than itself, equivalently, if it is generated only by itself and its reversal. For example, $abcd$ and $abcba$ are not primitive, because they are generated by $abcd$; but $abcdba$ is primitive. Note that factors of primitive words need not be primitive; for example, $abcdba$ is primitive, but its factor cb is not. Define a *primitive generator* of \bar{c} to be a generator of \bar{c} that is itself primitive. It follows from the foregoing remarks that every word \bar{c} has some primitive generator \bar{a} , and indeed, \bar{a} as well, since the reversal of a primitive generator is clearly a primitive generator. The following observation, on the other hand, is surprising.

LEMMA 2 (Thm. 1 of [30]). *The primitive generator of any word is unique up to reversal.*

For a very similar, though not identical, result, see [2].

Define the *primitive length* of any word \bar{c} to be the length of any primitive generator of \bar{c} . By Lemma 2, this notion is well-defined; it will play a significant role in our analysis of the adjacent fragment. Clearly, the primitive length of \bar{c} is at most $|\bar{c}|$, but will be strictly less if \bar{c} is not primitive.

It is important to realize that, while primitive generators are unique up to reversal, modes of generation are not. Indeed, $\bar{a} := abcba$ is one of the two primitive generators of $\bar{c} := abcba$, but we have $\bar{a}^f = \bar{c}$ for $f: [1, 7] \rightarrow [1, 5]$ given by either of the courses of values $[1, 2, 3, 4, 3, 4, 5]$ or $[1, 2, 3, 2, 3, 4, 5]$. In the sequel, it will be important to identify those words on which a given pair of surjective adjacent functions yield identical outputs. A *palindrome* is a word equal to its reversal; a palindrome is *non-trivial* if its length is at least 2. Let $\bar{a} := a_1 \cdots a_k$ be a word of length k . We say that a pair $\langle i, j \rangle$ is a *defect* of \bar{a} if the factor $a_i \cdots a_j$ is a non-trivial palindrome. We denote the set of defects of \bar{a} by $D_{\bar{a}}$, and regard it as a binary relation on the set $[1, k]$. If R is any binary relation, we write R^* for its equivalence closure, i.e. the smallest reflexive, symmetric and transitive relation that includes R . Now, for any pair of adjacent functions $f, g: [1, m] \rightarrow [1, k]$ and any set of pairs $D \subseteq \{\langle i, j \rangle \mid 1 \leq i < j \leq k\}$, we write $f \stackrel{D}{=} g$ if $\langle f(i), g(i) \rangle \in D^*$ for all i ($1 \leq i \leq m$). Evidently, $\stackrel{D}{=}$ is an equivalence relation.

LEMMA 3 (Thm. 4 of [30]). *Let \bar{a} be a primitive word of length k with defect set D , and let f and g be surjective functions in \mathbf{A}_k^m for some $m \geq k$. Then $\bar{a}^f = \bar{a}^g$ if and only if $f \stackrel{D}{=} g$.*

Of course, given surjective functions $f, g \in \mathbf{A}_k^m$ and any $D \subseteq \{\langle i, j \rangle \mid 1 \leq i < j \leq k\}$, it is a simple matter to check whether $f \stackrel{D}{=} g$. Lemma 3 allows us to read this condition as stating that f and g yield the same tuples when applied to any primitive word of length k whose defect set includes D .

Any adjacent function $f: [1, \ell], \rightarrow [1, k]$ induces a natural map from quantifier-free \mathcal{AF}^ℓ -formulas to quantifier-free \mathcal{AF}^k -formulas. Specifically, if χ is a quantifier-free \mathcal{AF}^ℓ -formula, denote by χ^g the formula $\chi(x_{g(1)} \cdots x_{g(\ell)})$, obtained by simultaneously replacing every variable x_i in χ by the corresponding variable $x_{g(i)}$. We claim that $\chi^g \in \mathcal{AF}^k$. Indeed, any atom α appearing in χ is of the form $p(\mathbf{x}_k^f)$, where p is a predicate of some arity m and $f \in \mathbf{A}_\ell^m$. But then the corresponding atom in χ^g has the form $\beta := \alpha(x_{g(1)} \cdots x_{g(\ell)}) = p(x_{g(f(1))} \cdots x_{g(f(m))}) = p(\mathbf{x}_k^{g \circ f})$. Since the composition of adjacent functions is adjacent, the assertion follows. The following (almost trivial) lemma is useful when manipulating adjacent formulas. Recall in this regard that any function $g \in \mathbf{A}_k^\ell$ maps a k -tuple \bar{a} over some set to an ℓ -tuple \bar{a}^g over the same set.

LEMMA 4. *Let χ be a quantifier-free \mathcal{AF}^ℓ -formula, and let $g \in \mathbf{A}_k^\ell$. For any structure \mathfrak{A} and any $\bar{a} \in A^k$, we have $\mathfrak{A} \models \chi^g[\bar{a}]$ if and only if $\mathfrak{A} \models \chi[\bar{a}^g]$. Thus, the adjacent type of any tuple in \mathfrak{A} is determined by that of its primitive generator.*

PROOF. We need consider only the case where χ is atomic: the general case follows by straightforward structural induction. Let $\chi := p(\mathbf{x}_k^f)$, with p an m -ary predicate, and $f \in \mathbf{A}_\ell^m$. Then, writing $\bar{a} = a_1 \cdots a_m$, both sides of the bi-conditional amount to the statement $a_{g(f(1))} \cdots a_{g(f(m))} \in p^{\mathfrak{A}}$. For the second statement, let \mathfrak{A} be a structure, and \bar{a} an ℓ -tuple from A . Then \bar{a} has a primitive generator, say \bar{b} of length $k \leq \ell$, with $\bar{a} = \bar{b}^g$ for some (surjective) $g \in \mathbf{A}_k^\ell$. Now consider any atomic \mathcal{AF}^ℓ -formula α . Then $\mathfrak{A} \models \alpha[\bar{a}]$ if and only if $\mathfrak{A} \models \alpha^g[\bar{b}]$. \dashv

Let \mathfrak{A} and \mathfrak{A}' be structures interpreting some common signature over a common domain A , and let $\ell \geq 0$. We write $\mathfrak{A} \approx_\ell \mathfrak{A}'$, if, for any predicate p of arity $m \geq 0$, and any m -tuple \bar{a} from A of primitive length at most ℓ , we have $\bar{a} \in p^{\mathfrak{A}}$ if and only if $\bar{a} \in p^{\mathfrak{A}'}$. That is, $\mathfrak{A} \approx_\ell \mathfrak{A}'$ just in case, for any predicate p interpreted by \mathfrak{A} , $p^{\mathfrak{A}}$ and $p^{\mathfrak{A}'}$ agree on all those m -tuples whose primitive length is at most ℓ . The next lemma states that, when evaluating \mathcal{AF}^ℓ -formulas in structures, we can disregard tuples whose primitive length is greater than ℓ .

LEMMA 5. *Let φ be an \mathcal{AF}^ℓ -sentence, and suppose \mathfrak{A} and \mathfrak{A}' are $\text{sig}(\varphi)$ -structures over a common domain A such that $\mathfrak{A} \approx_\ell \mathfrak{A}'$. Then $\mathfrak{A} \models \varphi \Leftrightarrow \mathfrak{A}' \models \varphi$.*

PROOF. Let ψ be a formula of \mathcal{AF}^ℓ (possibly featuring free variables), and let k (bounded by ℓ) be such that $\psi \in \mathcal{AF}^{[k]}$. We claim that, for any k -tuple of elements \bar{a} , $\mathfrak{A} \models \psi[\bar{a}]$ if and only if $\mathfrak{A}' \models \psi[\bar{a}]$. The statement of the lemma is the special case where ψ has no free variables. Again, we need consider only the case where ψ is atomic: the general case follows by straightforward structural

induction. Let $\psi := p(\mathbf{x}_k^f)$, with p an m -ary predicate, and $f \in \mathbf{A}_k^m$. If \bar{a} is a k -tuple of elements from A , then $\mathfrak{A} \models \psi[\bar{a}]$ if and only if $\bar{a}^f \in p^{\mathfrak{A}}$, and similarly for \mathfrak{A}' . But the primitive length of \bar{a}^f is certainly at most $k = |\bar{a}|$, and thus $\bar{a}^f \in p^{\mathfrak{A}}$ if and only if $\bar{a}^f \in p^{\mathfrak{A}'}$. \dashv

In view of Lemma 5, when considering models of \mathcal{AF}^ℓ -sentences, it will be useful to take the extensions of non-logical predicates (of whatever arity) to be undefined in respect of tuples whose primitive length is greater than ℓ , since these cannot affect the outcome of semantic evaluation. That is, where ℓ is clear from context, we typically suppose any model \mathfrak{A} of φ to determine whether $\bar{a} \in p^{\mathfrak{A}}$ for any m -ary predicate p and any m -tuple \bar{a} of primitive length at most ℓ ; but with respect to m -tuples \bar{a} having greater primitive length, \mathfrak{A} remains agnostic. To make it clear that the structure \mathfrak{A} need not be fully defined, we refer to it as a *layered structure*, of *height* ℓ . Notice that the notion of height is independent of the arities of the predicates interpreted. A layered structure \mathfrak{A} may have height, say 3, but still interpret a predicate p of arity, say, 5. In this case, it is determined whether $\mathfrak{A} \models p[babcc]$, because the primitive generator of $babcc$ is abc ; however, it is not determined whether $\mathfrak{A} \models p[abcab]$, because $abcab$ is primitive.

This idea enables us to build up models of \mathcal{AF} -formulas layer by layer. Suppose \mathfrak{A} is a layered structure of height ℓ , and we wish to construct a layered structure \mathfrak{A}^+ of height $\ell+1$ over the same domain A , agreeing with the assignments made by \mathfrak{A} . Clearly, it suffices to fix the adjacent type of each primitive $(\ell+1)$ -tuple \bar{b} from A . Suppose we want to fix the adjacent type of \bar{b} and hence that of its reversal \tilde{b} . To do so, we consider each predicate p in turn—of arity, say, m —and decide, for any m -tuple \bar{c} from A whose primitive generator is \bar{b} , whether $\mathfrak{A} \models p[\bar{c}]$. Now repeat this process for all pairs of mutually inverse primitive words (\bar{b}, \tilde{b}) from A having primitive length $\ell+1$. Since every tuple \bar{c} considered for inclusion in the extension of some predicate has primitive length $\ell+1$, these assignments will not clash with any previously made in the original structure \mathfrak{A} . Moreover, since, by Lemma 2, every m -tuple \bar{c} assigned in this process has a unique primitive generator \bar{b} (up to reversal), these assignments will not clash with each other. Thus, to elevate \mathfrak{A} to a layered structure of height $\ell+1$, one takes each inverse pair (\bar{b}, \tilde{b}) of primitive $(\ell+1)$ -tuples in turn, and fixes the adjacent type of each \bar{b} consistently with the existing assignments of all tuples generated by proper infixes of \bar{b} , as given in the original structure \mathfrak{A} .

We finish this section with an easy technical observation that will be needed in the sequel. Denote by $\vec{\mathbf{A}}_k^m$ the set of all functions $f \in \mathbf{A}_k^m$ such that $f(m) = k$. We refer to f as a *final adjacent function*. Thus, if $f \in \vec{\mathbf{A}}_k^m$ is thought of as a stroll of length m on some word \bar{a} of length k , then that stroll ends at the final position of \bar{a} .

LEMMA 6. *Let \bar{c} be a word of length m over some alphabet A , and d an element of A that does not appear in \bar{c} . If $\bar{c}d$ is not primitive, then neither is \bar{c} . In fact, there is a word \bar{a} of length $k < m$ and a function $f \in \vec{\mathbf{A}}_k^m$ such that $\bar{a}^f = \bar{c}$.*

PROOF. Suppose $\bar{c}d = \bar{b}^g$ for some word \bar{b} of length $k+1$ (bounded by m) and some surjective map $g \in \mathbf{A}_{k+1}^{m+1}$. Since d does not occur in \bar{c} , it is immediate that d occupies either the first or last position in \bar{b} , for otherwise, it would be

encountered again in the entire traversal of \bar{b} (as g is adjacent and surjective). By reversing \bar{b} if necessary, assume the latter, so that we may write $\bar{b} = \bar{a}d$, with $g(m+1) = k+1$. By adjacency, $g(m) = k$, so that setting $f = g \setminus \{ \langle m+1, k+1 \rangle \}$, we have the required \bar{a} and f . \dashv

We remark that, if $f \in \vec{\mathbf{A}}_k^m$, then the function $f^+ = f \cup \langle m+1, k+1 \rangle$ satisfies $f^+ \in \vec{\mathbf{A}}_{k+1}^{m+1}$. That is, we can extend f by setting $f(m+1) = k+1$, retaining adjacency. We utilise this fact as follows. Let φ be a normal-form $\mathcal{AF}^{\ell+1}$ -formula as given in (2). Recalling that we write $\chi^h = \chi(x_{h(1)} \cdots x_{h(\ell+1)})$ for χ quantifier-free in $\mathcal{AF}^{\ell+1}$ and $h \in \mathbf{A}_k^{\ell+1}$, we define the *adjacent closure* of φ , denoted $\text{acl}(\varphi)$, to be:

$$\bigwedge_{i \in I} \bigwedge_{k=1}^{\ell-1} \bigwedge_{f \in \vec{\mathbf{A}}_k^\ell} \forall \mathbf{x}_k \exists x_{k+1} (\gamma_i)^{f^+} \wedge \bigwedge_{k=1}^{\ell} \bigwedge_{g \in \mathbf{A}_k^{\ell+1}} \forall \mathbf{x}_k \beta^g.$$

Observe that the conjunctions for the $\forall \exists$ -formulas range over $f \in \vec{\mathbf{A}}_k^\ell$ (so that $f^+ \in \mathbf{A}^{\ell+1}$), while the conjunctions for the purely universal formula range over the whole of $\mathbf{A}_k^{\ell+1}$. Up to trivial logical rearrangement and re-indexing of variables, $\text{acl}(\varphi)$ is actually a normal-form \mathcal{AF}^ℓ -formula. In effect, $\text{acl}(\varphi)$ is the result of identifying various universally quantified variables in φ in a way which preserves adjacency. The following lemma is therefore immediate.

LEMMA 7. *Let $\varphi \in \mathcal{AF}^{\ell+1}$ be in normal-form. Then $\varphi \models \text{acl}(\varphi)$.*

Important notation mentioned in this and the previous sections is recapitulated in Table 1 for future reference.

§4. Upper bounds for \mathcal{AF}^ℓ without equality. Our goal in the following two sections is to establish a small model property for each of the fragments \mathcal{AF}^ℓ , for $\ell \geq 2$. We proceed by induction on ℓ . The base case ($\ell = 2$) involves no work: it was shown in [16, Thm. 5.1] that each satisfiable FO^2 -sentence φ has a model of size $2^{O(\|\varphi\|)}$; since $\mathcal{AF}^2 \subseteq \text{FO}^2$, the result follows. For the inductive step, we reduce the case $\ell+1$, with exponential blow-up, to the case ℓ . More precisely, we compute, for a given formula $\varphi \in \mathcal{AF}^{\ell+1}$, an equisatisfiable formula $\psi \in \mathcal{AF}^\ell$, over an exponentially larger signature, together with bounds on the relative sizes of their respective models. We thereby show that any satisfiable $\mathcal{AF}^{\ell+1}$ -sentence φ is satisfiable in a structure of size $\mathfrak{t}(\ell-1, O(\|\varphi\|))$, where \mathfrak{t} is defined inductively by $\mathfrak{t}(0, n) = n$ and $\mathfrak{t}(k+1, n) = 2^{\mathfrak{t}(k, n)}$. To illustrate the proof strategy as perspicuously as possible, we confine our attention in this section to the sub-fragment of \mathcal{AF}^ℓ *without equality*. In the next section we generalize the result to the full fragment \mathcal{AF}^ℓ , at the cost of some additional combinatorics. Returning to current affairs, Table 2 provides a summary of important symbols that will be defined throughout this section. (There is no need to read it just yet!)

For the next few lemmas (8–13), fix an equality-free, normal-form $\mathcal{AF}^{\ell+1}$ -formula φ over some signature τ , as given in (2), with $\ell \geq 2$. We construct an equality-free, normal-form formula $\psi \in \mathcal{AF}^\ell$ such that: (i) if φ is satisfiable over some domain A , then so is ψ ; and (ii) if ψ is satisfiable over a domain B ,

Functions $f : [1, m] \rightarrow [1, k]$ and tuples $\bar{a} = a_1 a_2 \cdots a_k$	
\mathbf{A}_k^m	the set of all adjacent functions $f : [1, m] \rightarrow [1, k]$
$\bar{\mathbf{A}}_k^m$	the set of all final adjacent functions $f : [1, m] \rightarrow [1, k]$
f^+	$f \cup \{(m+1, k+1)\}$ only if $f \in \bar{\mathbf{A}}_k^m$
\mathbf{x}_k	$x_1 x_2 \cdots x_k$
\bar{a}	reversal of \bar{a}
\bar{a}^f	$a_{f(1)} a_{f(2)} \cdots a_{f(m)}$
Formulas $\chi \in \mathcal{AF}^{[m]}$ with $f : [1, m] \rightarrow [1, k]$	
$\text{sig}(\chi)$	the signature of χ
χ^{-1}	$\chi(x_m x_{m-1} \cdots x_1)$
$\hat{\chi}$	$\chi \wedge \chi^{-1}$
χ^f	$\chi(\mathbf{x}_k^f)$
Normal form φ (in $(\ell+1)$ -variables)	
φ	$\bigwedge_{i \in I} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_i \wedge \forall \mathbf{x}_{\ell+1} \beta$
$\text{acl}(\varphi)$	$\bigwedge_{i \in I} \bigwedge_{k=1}^{\ell-1} \bigwedge_{f \in \bar{\mathbf{A}}_k^\ell} \forall \mathbf{x}_k \exists x_{k+1} \gamma_i^{f^+} \wedge \bigwedge_{k=1}^\ell \bigwedge_{g \in \mathbf{A}_k^{\ell+1}} \forall \mathbf{x}_k \beta^g$
Adjacent k -types χ and the k -type of a k -tuple \bar{a} in \mathfrak{A}	
Atp_k^τ	the set of all adjacent k -types over τ
$\text{atp}_k^{\mathfrak{A}}(\bar{a})$	the adjacent k -type of \bar{a} in \mathfrak{A}
$\text{itp}_k^{\mathfrak{A}}(\bar{a})$	the incremental k -type of \bar{a} in \mathfrak{A}
$\partial\xi$	the incremental k -type included in ξ
χ^+	$\chi(x_2 x_3 \cdots x_{k+1})$

TABLE 1. Quick reference guide for Sec 2 and 3.

then φ is satisfiable over some domain C , with $|C|/|B|$ bounded by an exponential function of $\|\varphi\|$. For the remainder of this section, all formulas will silently be assumed to be equality-free, and likewise for adjacent types.

We take ψ to have the form

$$\psi := \text{acl}(\varphi) \wedge \psi_1 \wedge \psi_2 \wedge \psi_3 \wedge \psi_4,$$

where $\text{acl}(\varphi)$ is the adjacent closure of φ (featured in Lemma 7) and ψ_1, \dots, ψ_4 are \mathcal{AF}^ℓ -formulas over an expanded signature. To motivate this construction of ψ , we first suppose φ is satisfiable, and consider any model $\mathfrak{A} \models \varphi$. We then introduce the conjuncts ψ_1, \dots, ψ_4 one by one, simultaneously defining an expansion \mathfrak{A}^+ of \mathfrak{A} by interpreting the new predicates so as to satisfy these conjuncts. (Of course, the ψ_1, \dots, ψ_4 depend only on φ , and not on \mathfrak{A}^+ .) Since, by Lemma 7, $\mathfrak{A} \models \text{acl}(\varphi)$, we have $\mathfrak{A}^+ \models \psi$. Our main task is then to show that, given any finite layered structure $\mathfrak{B} \models \psi$ of height ℓ , we can construct a finite layered structure $\mathfrak{C} \models \varphi$ of height $\ell+1$, with $|C|/|B|$ bounded by some exponential function of $\|\varphi\|$. This establishes the finite model property for $\mathcal{AF}^{\ell+1}$, and reduces its satisfiability problem to that for \mathcal{AF}^ℓ , though with exponential blow-up. In effect, ψ specifies just the right amount of information concerning the primitive ℓ -tuples occurring in any of its models to ensure that the adjacent

Surjective $f, g \in \mathbf{A}_k^m$, incremental k -types ι , quantifier-free $\chi \in \mathcal{AF}^{[k]}$	
\mathbf{D}_k^0	all pairs $\langle i, j \rangle$ for $1 \leq i < j \leq k$ with $j-i+1 \geq 2$ and odd
D^+	$\{\langle i+1, j+1 \rangle \mid \langle i, j \rangle \in D\}$
R^*	the equivalence closure of a binary relation R
$f \stackrel{D}{=} g$	$\langle f(i), g(i) \rangle \in D^*$ for all $i \in [1, m]$
ι is D -compatible	$f \stackrel{D}{=} g$ implies $\iota \models p(\mathbf{x}_k^f) \leftrightarrow p(\mathbf{x}_k^g)$ for all $p \in \text{sig}(\chi)$, all $f, g \in \mathbf{A}_k^m$
χ is D -consistent	there exists an $\xi \in \text{Atp}_k^\tau$ s.t. $\xi \models \chi$ and $\partial\xi$ is D -compatible
Construction of ψ from φ	
$d_k(\mathbf{x}_k)$	atom implying that \mathbf{x}_k is a palindrome
$\delta_D(\mathbf{x}_\ell)$	$\bigwedge \{d_{j-i+1}(x_i \cdots x_j) \mid \langle i, j \rangle \in D\}$
$p_\zeta(\mathbf{x}_{\ell-1})$	atom implying there is some x s.t. $x\mathbf{x}_{\ell-1}$ realizes ζ
$\mathfrak{B} \times H$	the model obtained by cloning elements of \mathfrak{B} for each $h \in H$

TABLE 2. Quick reference guide for Sec 4.

types of primitive $(\ell+1)$ -tuples can be assigned in such a way as to build a model of φ .

We now proceed to the definition of the conjuncts ψ_1, \dots, ψ_4 , and the construction of the expansion \mathfrak{A}^+ . Turning first to ψ_1 , for each s ($3 \leq 2s+1 \leq \ell$) we introduce a fresh $(2s+1)$ -ary predicate d_{2s+1} , and declare that a tuple $\bar{b} \in A^{2s+1}$ satisfies d_{2s+1} in \mathfrak{A}^+ just in case \bar{b} is a palindrome. It is then easy to verify that \mathfrak{A}^+ is a model of the sentence ψ_1 given as

$$\psi_1 := \bigwedge_{3 \leq 2s+1 \leq \ell} \forall \mathbf{x}_{k+1} d_{2s+1}(x_1 \cdots x_s x_{s+1} x_s \cdots x_1).$$

Conversely, if \mathfrak{B} is any structure such that $\mathfrak{B} \models \psi_1$, and $\bar{c} \in B^{2s+1}$ is a palindrome ($3 \leq 2s+1 \leq \ell$), then $\mathfrak{B} \models d_{2s+1}[\bar{c}]$. Observe that ψ_1 employs at most $(\ell+1)/2 < \ell$ variables.

The predicates d_s just introduced will be useful for specifying that various tuples of objects exhibit certain sets of *defects*, in the sense of Sec. 3. Recall that a defect of an ℓ -tuple $a_1 \cdots a_\ell$ is a pair of integers $\langle i, j \rangle$ such that $a_i \cdots a_j$ is a non-trivial palindrome. (Note that the length of this factor is $j-i+1 \geq 2$.) In the sequel, for any $k \geq 2$, we denote by \mathbf{D}_k^0 the set of all pairs $\langle i, j \rangle$ for $1 \leq i < j \leq k$ such that $j-i+1$ is greater than two and odd. (Thus, $\mathbf{D}_2^0 = \emptyset$.) Now, for any $D \subseteq \mathbf{D}_{\ell-1}^0$, we write δ_D for the formula

$$\delta_D := \bigwedge \{d_{j-i+1}(x_i \cdots x_j) \mid \langle i, j \rangle \in D\}.$$

Intuitively, δ_D says that any satisfying tuple has a defect set which includes D . Note that δ_D is defined only for $D \subseteq \mathbf{D}_{\ell-1}^0$: for technical reasons, we are not interested in defects corresponding to even-length palindromic factors.

Turning now to ψ_2 , we introduce, for each adjacent ℓ -type $\zeta \in \text{Atp}_\ell^\tau$, an $(\ell-1)$ -ary predicate p_ζ intended to identify the tails of ℓ -tuples satisfying ζ . Specifically, we declare that $\mathfrak{A}^+ \models p_\zeta[\bar{b}]$ just in case there is some $a \in A$ such that $\mathfrak{A} \models \zeta[a\bar{b}]$.

It is then easy to verify that \mathfrak{A}^+ is a model of the sentence ψ_2 given as

$$\psi_2 := \bigwedge_{\zeta \in \text{Atp}_\ell^\tau} \forall \mathbf{x}_\ell \left(\zeta \rightarrow p_\zeta(x_2 \cdots x_\ell) \right).$$

Conversely, if \mathfrak{B} is a structure interpreting the signature τ , and \mathfrak{B}^+ is an expansion of \mathfrak{B} such that $\mathfrak{B}^+ \models \psi_2$, then $\text{atp}^{\mathfrak{B}}[a\bar{b}] = \zeta$ implies $\mathfrak{B}^+ \models p_\zeta[\bar{b}]$. Observe that ψ_2 employs ℓ variables.

The construction of ψ_3 and ψ_4 requires some preliminary work concerning palindromes, as tuples containing palindromic factors pose a technical challenge in the construction of our formula ψ . To explain why, let \bar{a} be an $(\ell+1)$ -tuple over some set A and ξ an adjacent $(\ell+1)$ -type, and imagine that we wish to define a structure \mathfrak{A} on A in such a way that $\text{atp}^{\mathfrak{A}}[\bar{a}] = \xi$. Is this at all possible? Suppose, for example, that $\bar{a} = abcba$, and that ξ contains the literals $p(x_1x_2x_3x_2x_3x_4x_5)$ and $\neg p(x_1x_2x_3x_4x_3x_4x_5)$. A moment's thought shows that we cannot have $\text{atp}^{\mathfrak{A}}[\bar{a}] = \xi$, since that would require both $\mathfrak{A} \models p[abcba]$ and $\mathfrak{A} \models \neg p[abcba]$. We need to be able to identify this sort of situation using only the resources of \mathcal{AF} . The next lemma explains how.

Recall the apparatus of incremental types introduced in Sec. 2: if \bar{c} is a k -tuple over some structure \mathfrak{A} , then $\text{itp}^{\mathfrak{A}}[\bar{c}]$ is the set of covering adjacent k -literals (i.e. those featuring all variables in \mathbf{x}_k) satisfied by \bar{c} in \mathfrak{A} . Recall also the notation introduced in Lemma 3: for any pair of adjacent functions $f, g: [1, m] \rightarrow [1, k]$ and any set $D \subseteq \{\langle i, j \rangle \mid 1 \leq i < j \leq k\}$, we write $f \stackrel{D}{=} g$ if $\langle f(i), g(i) \rangle \in D^*$ for all i ($1 \leq i \leq m$), where D^* denotes the equivalence closure of D . Let us say that an incremental k -type ι is D -compatible if $f \stackrel{D}{=} g$ implies $\iota \models p(\mathbf{x}_k^f) \leftrightarrow p(\mathbf{x}_k^g)$ for all predicates p in the signature of ι and all surjective adjacent functions $f, g: [1, m] \rightarrow [1, k]$ where m is the arity of p .

LEMMA 8. *Let \bar{c} be a primitive k -tuple over A , and D the defect set of \bar{c} . If \mathfrak{A} is a structure with domain A , then $\iota = \text{itp}^{\mathfrak{A}}[\bar{c}]$ is D -compatible. Conversely, if ι is a D -compatible incremental k -type over some signature τ , then there is a structure \mathfrak{A} interpreting τ over A such that $\text{itp}^{\mathfrak{A}}[\bar{c}] = \iota$.*

PROOF. The first statement of the lemma is almost immediate. Fix a predicate p of arity m interpreted by \mathfrak{A} , and suppose $f, g \in \mathbf{A}_k^m$ are surjective. We must show that $f \stackrel{D}{=} g$ implies $\iota \models p(\mathbf{x}_k^f) \leftrightarrow p(\mathbf{x}_k^g)$. But by Lemma 3, if $f \stackrel{D}{=} g$, then $\bar{c}^f = \bar{c}^g$, hence $\mathfrak{A} \models p[\bar{c}^f] \leftrightarrow \mathfrak{A} \models p[\bar{c}^g]$. Thus $\iota \models p(\mathbf{x}_k^f) \leftrightarrow p(\mathbf{x}_k^g)$.

For the second statement, define \mathfrak{A} over the domain A by setting, for any predicate p of τ having arity, say, m :

$$p^{\mathfrak{A}} = \{\bar{a}^f \mid \text{there exists a surjective } f \in \mathbf{A}_k^m \text{ such that } \iota \models p(\mathbf{x}_k^f)\}.$$

To show that $\mathfrak{A} \models \iota[\bar{a}]$, fix any $p \in \tau$ with arity m . If ι contains the (covering, adjacent) atom $\alpha := p(\bar{x}^f)$, where then it is immediate from the construction of \mathfrak{A} that $\mathfrak{A} \models \alpha[\bar{a}]$. It remains to show that if ι contains the negated atom $\nu := \neg p(\mathbf{x}_k^f)$, then $\mathfrak{A} \models \nu[\bar{a}]$. Suppose otherwise. From the construction of \mathfrak{A} , we have $\iota \models p(\mathbf{x}_k^g)$ for some surjective adjacent $g: [1, m] \rightarrow [1, k]$ such that $\bar{a}^f = \bar{a}^g$. By Lemma 3, $f \stackrel{D}{=} g$. Yet ι is by assumption D -compatible, whence

$\iota \models p(\bar{x}^f) \leftrightarrow p(\bar{x}^g)$, contradicting the fact that ξ contains both $\neg p(\mathbf{x}_k^f)$ and $p(\mathbf{x}_k^g)$. \dashv

To explain the significance of this lemma, let us return to our example of the 5-tuple $\bar{a} = abcbd$, and the adjacent 5-type ξ containing literals $p(x_1x_2x_3x_2x_3x_4x_5)$ and $\neg p(x_1x_2x_3x_4x_3x_4x_5)$. Notice that these literals are covering, and thus are contained in $\partial\xi$, the incremental 5-type included in ξ . As we have seen, \bar{a} cannot be assigned the adjacent-type ξ , because it makes inconsistent demands in respect of the 7-tuple $abcbcbd$. Observe, however, that the defect set of \bar{a} is $D = \{2, 4\}$, and when we write the argument sequences $x_1x_2x_3x_2x_3x_4x_5$ and $x_1x_2x_3x_4x_3x_4x_5$ in the form \mathbf{x}_5^f and \mathbf{x}_5^g , respectively, it is easily checked that $f \stackrel{D}{=} g$, whence $\partial\xi$ is not D -compatible. On the other hand, setting $\ell := 4$, and supposing the formula ψ_1 to hold, we see that $bc b$ satisfies the predicate d_3 , whence the quadruple $abcb$ satisfies the \mathcal{AF}^4 -formula $\delta_D = d_3(x_2x_3x_4)$. Crucially, the fact that $abcb$ satisfies δ_D allows us to detect, using resources available only in \mathcal{AF}^4 , that $\bar{a} = abcbd$ cannot be assigned the adjacent 5-type ξ , because $\partial\xi$ is not D -compatible.

Returning now to the construction of ψ_3 and ψ_4 , and recalling that $\mathbf{D}_{\ell-1}^\circ$ denotes the set of all pairs $\langle i, j \rangle$ for $1 \leq i < j \leq \ell-1$ such that $j-i+1$ is greater than 2 and odd, consider any subset $D \subseteq \mathbf{D}_{\ell-1}^\circ$. Denote by D^+ the result of adding 1 to both components of every pair in D , i.e. $D^+ := \{\langle i+1, j+1 \rangle \mid \langle i, j \rangle \in D\}$. (Thus, $D^+ \subseteq \mathbf{D}_\ell^\circ \subseteq \mathbf{D}_{\ell+1}^\circ$.) The intuition here is that if D represents a set of odd-length non-trivial palindromic factors of an $(\ell-1)$ -tuple \bar{a} , and a and b are some elements, then D^+ represents a set of odd-length non-trivial palindromic factors of the ℓ -tuple $a\bar{a}$ and indeed also of the $(\ell+1)$ -tuple $a\bar{a}b$. Say that a quantifier-free $\mathcal{AF}^{\ell+1}$ -formula χ is D^+ -consistent if there exists an adjacent $(\ell+1)$ -type ξ over the relevant signature such that $\xi \models \chi$ and $\partial\xi$ is D^+ -compatible. Finally, define the formulas ψ_3 and ψ_4 as:

$$\begin{aligned} \psi_3 &:= \bigwedge_{i \in I} \bigwedge_{\zeta \in \text{Atp}_\ell^\tau} \bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^\circ} \forall \mathbf{x}_{\ell-1} \exists x_\ell \left((\delta_D \wedge p_\zeta(\mathbf{x}_{\ell-1})) \rightarrow \right. \\ &\quad \left. \bigvee \{ \eta \in \text{Atp}_\ell^\tau \mid (\zeta \wedge \hat{\beta} \wedge \gamma_i \wedge \eta^+) \text{ is } D^+\text{-consistent} \} \right) \\ \psi_4 &:= \bigwedge_{\zeta \in \text{Atp}_\ell^\tau} \bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^\circ} \forall \mathbf{x}_\ell \left((\delta_D \wedge p_\zeta(\mathbf{x}_{\ell-1})) \rightarrow \right. \\ &\quad \left. \bigvee \{ \eta \in \text{Atp}_\ell^\tau \mid (\zeta \wedge \hat{\beta} \wedge \eta^+) \text{ is } D^+\text{-consistent} \} \right). \end{aligned}$$

We claim that $\mathfrak{A}^+ \models \psi_3$. Suppose \bar{a} is an $(\ell-1)$ -tuple such that $\mathfrak{A}^+ \models \delta_D[\bar{a}]$ and $\mathfrak{A}^+ \models p_\zeta[\bar{a}]$. By construction of \mathfrak{A}^+ , there exists $a \in A$ such that $\text{atp}^{\mathfrak{A}^+}[a\bar{a}] = \zeta$. Now, fix an index $i \in I$ of some witness requirement in φ . Since $\mathfrak{A} \models \varphi$, there exists an element $b \in A$ such that $\mathfrak{A} \models \gamma_i[a\bar{a}b]$, and in fact such that $\mathfrak{A} \models (\hat{\beta} \wedge \gamma_i)[a\bar{a}b]$. Letting $\xi := \text{atp}^{\mathfrak{A}^+}[a\bar{a}b]$ and $\eta := \text{atp}^{\mathfrak{A}^+}[\bar{a}b]$, we have $\xi \models \zeta \wedge \hat{\beta} \wedge \gamma_i \wedge \eta^+$. On the other hand, since $\mathfrak{A}^+ \models \delta_D[\bar{a}]$ we have, by construction of \mathfrak{A}^+ , that the defect set of \bar{a} includes D ; and hence that the defect set of $a\bar{a}b$ includes D^+ . By the first statement of Lemma 8 we have that $\partial\xi$ is D^+ -compatible. Hence, b is a witness for the quantifier $\exists x_\ell$ in ψ_3 with respect to \bar{a} , so that $\mathfrak{A}^+ \models \psi_3$ as required. By similar reasoning, $\mathfrak{A}^+ \models \psi_4$.

To motivate these formulas, consider first ψ_3 , and suppose $\psi_1 \wedge \psi_2 \wedge \psi_3$ has a model, say, \mathfrak{B} . We may assume, by Lemma 5, that \mathfrak{B} is a layered structure of height ℓ , and we may take \mathfrak{B}^- to be the reduct of \mathfrak{B} to the signature τ (i.e. we forget the predicates p_ζ and d_{2s+1}). Suppose, now a is an element of B and \bar{a} an $(\ell-1)$ -tuple over B , and let $\zeta = \text{atp}^{\mathfrak{B}^-}[a\bar{a}]$. Let us further imagine that the defect set of \bar{a} is $D \subseteq \mathbf{D}_{\ell-1}^o$. (We shall not worry about even-length palindromic factors for the present.) Since $\mathfrak{B} \models \psi_1 \wedge \psi_2$, the $(\ell-1)$ -tuple \bar{a} will satisfy δ_D and p_ζ in \mathfrak{B} . Of course, \mathfrak{B} does not in general define the adjacent types of $(\ell+1)$ -tuples (unless they are non-primitive), but we would like to be able to find, for each index $i \in I$, an element $b \in B$ such that the $(\ell+1)$ -tuple $a\bar{a}b$ could be assigned an adjacent type satisfying the existential requirement γ_i of φ without violating the universal constraints β . Anticipating Lemma 13, this is what ψ_3 does: the witnesses it provides for x_ℓ have the required properties. Formula ψ_4 plays an analogous role with respect to the universal requirements of φ . But we are getting ahead of ourselves. For the present, we have proved:

LEMMA 9. *Suppose $\mathfrak{A} \models \varphi$. Then we can expand \mathfrak{A} to a model $\mathfrak{A}^+ \models \psi$.*

Having defined ψ and established Lemma 9, we wish to argue for the converse direction. We begin with three straightforward, technical lemmas. When dealing with equality-free formulas, we may freely duplicate elements in their models. Let \mathfrak{B} be a τ -structure, and H a non-empty set of indices. We define the structure $\mathfrak{B} \times H$ over the Cartesian product $B \times H$ as follows: for any $p \in \tau$ of arity k , and any k -tuples $b_1 \cdots b_k$ over B and $h_1 \cdots h_k$ over H , set $\mathfrak{B} \times H \models p[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$ if and only if $\mathfrak{B} \models p[b_1 \cdots b_k]$. By routine structural induction, we have:

LEMMA 10. *Let ψ be an equality-free first-order formula all of whose free variables occur in \mathbf{x}_k , \mathfrak{B} a structure interpreting the signature of ψ , and H a non-empty set. Then, for any tuples $\bar{b} = b_1 \cdots b_k$ from B and $h_1 \cdots h_k$ from H , $\mathfrak{B} \models \psi[\bar{b}]$ if and only if $\mathfrak{B} \times H \models \psi[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$.*

The following very simple lemma will be used in conjunction with Lemma 10.

LEMMA 11. *Suppose \bar{b} is a k -tuple over some set B and H a non-empty set. Then, for any $h_1, \dots, h_k \in H$, if \bar{c} is the tuple $\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle$ over $B \times H$, then every defect of \bar{c} is a defect of \bar{b} .*

PROOF. Pick some defect $\langle i, j \rangle$ of \bar{c} . Thus, $c_i \cdots c_j$ is a non-trivial palindrome, i.e. $\langle b_{i+m}, k_{i+m} \rangle = \langle b_{j-m}, k_{j-m} \rangle$ for all $0 \leq m \leq \lfloor (j-i)/2 \rfloor$. But then $b_i \cdots b_j$ is certainly a non-trivial palindrome, and hence $\langle i, j \rangle$ a defect of \bar{b} as required. \dashv

Finally, the following combinatorial lemma allows us to extend the technique of ‘circular witnessing’ (see, e.g. [16, p. 62–64]) frequently used in the analysis of two-variable logics, to the languages \mathcal{AF}^k .

LEMMA 12. *For any integer $k > 0$ there is a set H with $|H| = (k^2 + k + 1)^{k+1}$ and a function $g: H^k \rightarrow H$ such that, for any tuple $\bar{t} \in H^k$ consisting of the elements t_1, \dots, t_k in some order: (i) $g(\bar{t})$ is not in \bar{t} ; (ii) if $\bar{t}' \in H^k$ consists of the elements $\{t_2, \dots, t_k, g(\bar{t})\}$ in some order, then $g(\bar{t}')$ is not in \bar{t}' either.*

PROOF. Let $z = k^2 + k + 1$ and $H = [1, z]^{k+1}$. Thus, $|H|$ is as required. Writing any element of H as the word $i\bar{s}$, where $i \in [1, z]$ and $\bar{s} \in [1, z]^k$, let g be defined by

$g(i_1\bar{s}_1, \dots, i_k\bar{s}_k) = i_0i_1 \cdots i_k$, where i_0 is the smallest positive integer not in the set $S = \{i_1, \dots, i_k\} \cup \bar{s}_1 \cup \cdots \cup \bar{s}_k$. (For brevity, we are here identifying words over the integers with the sets of their members.) Note that $i_0 \in [1, z]$ since $|S| < z$. Now let some tuple $\bar{t} \in H^k$ be given, consisting of words t_1, \dots, t_k in some order, and write $t_h = i_h\bar{s}_h$ (for all $1 \leq h \leq k$), where $i_h \in [1, z]$ and $\bar{s}_h \in [1, z]^k$. Thus, $t = g(\bar{t})$ is a word of the form $i_0\bar{s}$, where \bar{s} consists of i_1, \dots, i_k in some order, and i_0 does not occur anywhere in \bar{t} . Condition (i) is then immediate because of the choice of i_0 . For condition (ii), we observe that $t' = g(\bar{t}')$ is a word of the form $i'\bar{s}'$, where \bar{s}' consists of i_0, i_2, \dots, i_k in some order, and i' does not occur in any of the words in \bar{t}' . By the choice of i' , it is immediate that $t' \notin \{t_2, \dots, t_k\}$. But the value i_1 (the first letter of t_1) occurs in $g(\bar{t})$ (which belongs to the tuple \bar{t}'), whence $i' \neq i_1$. It follows that $t' \neq t_1$, whence t' is not in \bar{t} , as required. \dashv

With the help of these three technical lemmas, we provide the advertised converse to Lemma 9.

LEMMA 13. *Suppose $\mathfrak{B} \models \psi$. Then we can construct a model $\mathfrak{C}^+ \models \varphi$ such that $|C^+|/|B| \leq |I| \cdot (\ell^2 + \ell + 1)^\ell$.*

PROOF. We repeat here for convenience the formula φ as defined in (2):

$$\bigwedge_{i \in I} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_i \wedge \forall \mathbf{x}_{\ell+1} \beta.$$

In the sequel, we refer to the conjunct $\forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_i$ as the *i th witness requirement* of φ . Our first step is to define a collection of ‘pseudo-witnesses’ in the structure \mathfrak{B} for the various witness requirements of φ . These will be used later to select actual witnesses in the model \mathfrak{C}^+ of φ that we eventually construct. Let \mathfrak{B}^- be the reduct of \mathfrak{B} to the signature τ (i.e. we forget the predicates p_ζ and d_{2s+1}). Consider any ℓ -tuple $\bar{b} = b_1 \cdots b_\ell$: let $\zeta = \text{atp}^{\mathfrak{B}^-}[\bar{b}]$ and let D be the set of defects of $b_2 \cdots b_\ell$ corresponding to odd-length palindromes, i.e., those defects $\langle i, j \rangle$ for which $j - i + 1$ is odd. Since $\mathfrak{B} \models \psi_1 \wedge \psi_2$, we have $\mathfrak{B} \models \delta_D[b_2 \cdots b_\ell]$ and $\mathfrak{B} \models p_\zeta[b_2 \cdots b_\ell]$. And since $\mathfrak{B} \models \psi_3$, given any $i \in I$, we may select $b \in B$ such that $\eta = \text{atp}^{\mathfrak{B}}[b_2 \cdots b_\ell b]$ is one of the disjuncts in the consequent

$$\bigvee \{ \eta \in \text{Atp}_\ell^\tau \mid (\zeta \wedge \hat{\beta} \wedge \gamma_i \wedge \eta^+) \text{ is } D^+\text{-consistent} \}$$

of the implication occurring in ψ_3 . It follows that there exists some $\xi \in \text{Atp}_{\ell+1}^\tau$ such that $\xi \models (\zeta \wedge \hat{\beta} \wedge \gamma_i \wedge \eta^+)$ and $\partial\xi$ is D^+ -compatible. We identify b as the ‘pseudo-witness’ for the i th witness requirement of φ , and write $b_{\bar{b}, i} := b$ and $\xi_{\bar{b}, i} := \xi$ to make the dependencies on i and \bar{b} explicit.

Next we inflate the model \mathfrak{B} using the construction of Lemma 10. Let H be a set of cardinality $\ell^2 + \ell + 1$ and let $g: H^\ell \rightarrow H$ be a function satisfying conditions (i) and (ii) guaranteed by Lemma 12. Define $\mathfrak{C} = \mathfrak{B}^- \times (I \times H)$, and define the layered structure \mathfrak{C}' (with the same domain, C) to be the result of disregarding any tuples in \mathfrak{C} of primitive length greater than ℓ . Thus, we may write elements of C as triples (b, i, j) , where $b \in B$, $i \in I$ and $j \in H$; and of course, \mathfrak{C}' has height ℓ . By Lemma 10, $\mathfrak{C} \models \psi$, and bearing in mind that $\psi \in \mathcal{AF}^\ell$, by Lemma 5, $\mathfrak{C}' \models \psi$. We may now define a collection of ‘witness functions’ $w_i: C^\ell \rightarrow C$, where i ranges over I . Consider any ℓ -tuple $\bar{c} := c_1 \cdots c_\ell$ of elements in C , with

$c_h = (b_h, i_h, j_h)$ for each h ($1 \leq h \leq \ell$). Writing $\bar{b} = b_1 \cdots b_\ell$, and recalling the pseudo-witness $b_{\bar{b},i}$ identified in the previous paragraph, we define $w_i(\bar{c})$ to be the element $(b_{\bar{b},i}, i, g(j_1 \cdots j_\ell))$. Thus, $w_i(\bar{c})$ is a judiciously chosen copy of $b_{\bar{b},i}$, able, potentially, to serve as an i th witness for φ . Most of the work will be done by its first component, $b_{\bar{b},i}$; the remaining components, i and $g(j_1 \cdots j_\ell)$, simply ensure that the eventually selected witnesses do not, as it were, tread on each others' toes. Note that the functions w_i satisfy the following properties:

- (w1): for fixed \bar{c} , the $w_i(\bar{c})$ are distinct as i varies over I ;
- (w2): $w_i(\bar{c})$ does not occur in \bar{c} ;
- (w3): if $\bar{c} = c_1 \cdots c_\ell$ and \bar{c}' is an ℓ -tuple consisting of the elements c_2, \dots, c_ℓ , $w_i(\bar{c})$ in some order, then $w_{i'}(\bar{c}')$ does not occur in \bar{c} for any $i' \in I$.

Indeed, (w1) is immediate from the fact $w_i(\bar{c})$ contains i as its second element; (w2) and (w3) follow, respectively, from conditions (i) and (ii) on g guaranteed in Lemma 12.

We are now ready to elevate the layered structure \mathfrak{C}' to a layered structure \mathfrak{C}^+ of height $\ell+1$ such that $\mathfrak{C}^+ \models \varphi$. We first ensure that \mathfrak{C}^+ provides witnesses for the various witness requirements of φ . Fix any ℓ -tuple $\bar{c} = c_1 \cdots c_\ell$ and any $i \in I$. We have two cases, depending on whether the $(\ell+1)$ -tuple $\bar{c} w_i(\bar{c})$ is primitive. Suppose first that it is not. By (w2), $w_i(\bar{c})$ is not an element of \bar{c} , whence by Lemma 6 there is some k -tuple \bar{d} ($k < \ell$) and $f \in \vec{\mathbf{A}}_k^\ell$ such that $\bar{d} = \bar{c}^f$. As before, define $f^+ \in \vec{\mathbf{A}}_{k+1}^{\ell+1}$ extending f by setting $f^+(k+1) = \ell+1$. Since $k < \ell$, and $\mathfrak{C}' \models \text{acl}(\varphi)$, there exists $c' \in C$ such that $\mathfrak{C}' \models \gamma_i^{f^+}[\bar{d}c']$. By Lemma 4, $\mathfrak{C}' \models \gamma_i[(\bar{d}c')^{f^+}]$, or in other words, $\mathfrak{C}' \models \gamma_i[\bar{c}c']$, so that a witness c' is already present in respect of the tuple \bar{c} and the index i . (Notice that we are throwing the element $w_i(\bar{c})$ away, as an i th witness is already present for the tuple \bar{c} .) Suppose on the other hand that $\bar{c} w_i(\bar{c})$ is primitive. Recall that, for all h ($1 \leq h \leq \ell$), the element c_h has the form (b_h, i_h, j_h) , and that $w_i(\bar{c}) = (b_{\bar{b},i}, i, g(j_1 \cdots j_\ell))$. Recall in addition the adjacent $(\ell+1)$ -type $\xi_{\bar{b},i}$ identified earlier. To reduce notational clutter, we write c for $w_i(\bar{c})$ and ξ for $\xi_{\bar{b},i}$. Setting $\zeta = \text{atp}^{\mathfrak{B}^-}[\bar{b}]$ and $\eta = \text{atp}^{\mathfrak{B}^-}[b_2 \cdots b_\ell b_{\bar{b},i}]$, Lemma 10 guarantees that $\zeta = \text{atp}^{\mathfrak{C}'}[\bar{c}]$ and $\eta = \text{atp}^{\mathfrak{C}'}[c_2 \cdots c_\ell c]$. Moreover, by Lemma 11, any defect of $c_2 \cdots c_\ell$ is evidently a defect of $b_2 \cdots b_\ell$. In fact, since \bar{c} is primitive and $c = w_i(\bar{c})$ does not occur in \bar{c} , the tuple $\bar{c}c$ cannot have any defects corresponding to even-length palindromes, and cannot have any defects of the forms $\langle 1, j \rangle$ or $\langle j, \ell+1 \rangle$. Thus, letting D denote the set of defects of $b_2 \cdots b_\ell$ corresponding to odd-length palindromes, we see that the defect set of $\bar{c}c$ is D^+ . But we have already argued that $\xi \models (\zeta \wedge \hat{\beta} \wedge \gamma_i \wedge \eta^+)$, and moreover that $\partial\xi$ is D^+ -compatible. By Lemma 8, there is a structure over the elements of $\bar{c}c$ in which the incremental type of $\bar{c}c$ is $\partial\xi$. Thus we may consistently fix $\text{itp}^{\mathfrak{C}^+}[\bar{c}c]$ to be $\partial\xi$. (Observe that this also fixes the incremental type of the reversed tuple $c\bar{c}$.) Since $\zeta = \text{atp}^{\mathfrak{C}'}[\bar{c}]$, $\eta = \text{atp}^{\mathfrak{C}'}[c_2 \cdots c_\ell c]$, and $\xi = \zeta \cup \eta^+ \cup \partial\xi$, we have $\text{atp}^{\mathfrak{C}^+}[\bar{c}c] = \xi$. Moreover, since $\xi \models \gamma_i$, c is an i th witness for φ with respect to the tuple \bar{c} . Finally, since $\xi \models \hat{\beta}$ the adjacent $(\ell+1)$ -types assigned to the tuples $\bar{c}c$ and $c\bar{c}$ satisfy β . Still keeping \bar{c} fixed for the moment, we may carry out the above procedure for all $i \in I$. To see that these assignments do not interfere with each other, we invoke property (w1) of the functions w_i .

Now make these assignments as just described for *each* ℓ -tuple \bar{c} over C . To ensure that these assignments do not interfere with each other, we make use of properties (w1) and (w3) of the functions w_i . If \bar{d} is an m -tuple that has been assigned (or not) to the extensions of various predicates by the process described above, then the two primitive generators of \bar{d} must be of the form $\bar{c}\bar{c}$ and $c\bar{c}$, where $c = w_i(\bar{c})$ for some $i \in I$. Since primitive generators are unique up to reversal by Theorem 2, it suffices to show that, for distinct pairs (\bar{c}, i) and (\bar{c}', i') , the corresponding $(\ell+1)$ -tuples $\bar{c}w_i(\bar{c})$ and $\bar{c}'w_{i'}(\bar{c}')$ are not the same up to reversal. Now $\bar{c}w_i(\bar{c}) = \bar{c}'w_{i'}(\bar{c}')$ implies $\bar{c} = \bar{c}'$, whence i and i' are distinct, whence $w_i(\bar{c}) \neq w_{i'}(\bar{c}')$ by (w1), a contradiction. On the other hand $\bar{c}'w_{i'}(\bar{c}') = w_i(\bar{c})\bar{c}$ implies $\bar{c}' = w_i(\bar{c}), c_\ell \cdots c_2$, whence $w_{i'}(\bar{c}')$ does not occur in \bar{c} (or therefore in \bar{c}) by (w3), again a contradiction.

At this point, we have assigned a collection of tuples with primitive length $\ell+1$ to the extensions of predicates in τ so as to guarantee that $\mathfrak{C}^+ \models \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_i$ for all $i \in I$. In addition, no adjacent $(\ell+1)$ -types thus defined violate $\forall \mathbf{x}_{\ell+1} \beta$. It remains to complete the specification of \mathfrak{C}^+ by defining the adjacent types of all remaining primitive $\ell+1$ -tuples, and showing that, in the resulting structure, every $(\ell+1)$ -tuple (primitive or not) satisfies β . Let $\bar{c} = c_1 \cdots c_{\ell+1}$ be a primitive $(\ell+1)$ -tuple whose adjacent type has not yet been defined. Let $\zeta = \text{atp}^{\mathfrak{C}}[c_1 \cdots c_\ell]$ and $\eta = \text{atp}^{\mathfrak{C}}[c_2 \cdots c_{\ell+1}]$. Since \bar{c} is primitive, it has no defects corresponding to even-length palindromes, and indeed no defects of the forms $\langle 1, j \rangle$ or $\langle j, \ell+1 \rangle$. Letting D be the defect set of $c_2 \cdots c_\ell$, therefore, we have $D \subseteq \mathbf{D}_{\ell-1}^0$; moreover, the defect set of \bar{c} is D^+ . Further, since $\mathfrak{B} \models \psi_1 \wedge \psi_2$, we have $\mathfrak{B} \models \delta_D[c_2 \cdots c_\ell]$ and $\mathfrak{B} \models p_\zeta[c_2 \cdots c_\ell]$. And since $\mathfrak{B} \models \psi_4$, we see that $\eta = \text{atp}^{\mathfrak{C}'}[c_2 \cdots c_{\ell+1}]$ is one of the disjuncts in the consequent of ψ_4 . Thus, there exists an adjacent $(\ell+1)$ -type ξ such that $\xi \models (\zeta \wedge \hat{\beta} \wedge \eta^+)$ and $\partial\xi$ is D^+ -compatible. By Lemma 8, therefore, there is a structure over the elements of \bar{c} in which \bar{c} has the incremental type $\partial\xi$. Thus we may consistently fix $\text{itp}^{\mathfrak{C}^+}[\bar{c}]$ to be $\partial\xi$. Since $\zeta = \text{atp}^{\mathfrak{C}'}[c_1 \cdots c_\ell]$, $\eta = \text{atp}^{\mathfrak{C}'}[c_2 \cdots c_{\ell+1}]$, and $\xi = \zeta \cup \eta^+ \cup \partial\xi$, we have $\text{atp}^{\mathfrak{C}^+}[\bar{c}] = \xi$. And since $\xi \models \hat{\beta}$, the adjacent $(\ell+1)$ -types assigned to the tuples \bar{c} and \bar{c} satisfy β . Carrying this procedure out for all remaining primitive $(\ell+1)$ -tuples, we obtain a layered structure \mathfrak{C}^+ of depth $\ell+1$. Let \bar{d} be any $(\ell+1)$ -tuple of elements from C . If \bar{d} is primitive, then we have just ensured that $\mathfrak{C}^+ \models \beta[\bar{d}]$. If, on the other hand, $\bar{d} = \bar{e}^f$ for some k -tuple \bar{e} and some $f \in \mathbf{A}_k^{\ell+1}$, where $k \leq \ell$, then, since $\mathfrak{C} \models \text{acl}(\varphi)$, we have $\mathfrak{C} \models \beta^f[\bar{e}]$ and hence, by Lemma 4, $\mathfrak{C} \models \beta[\bar{d}]$. This completes the construction of \mathfrak{C}^+ . We have shown that $\mathfrak{C}^+ \models \varphi$. \dashv

Let us take stock. Since \mathcal{AF}^2 is included in FO^2 , we know that each satisfiable \mathcal{AF}^2 -sentence φ has a finite model of size $2^{O(\|\varphi\|)}$. Further, given any equality-free normal-form $\mathcal{AF}^{\ell+1}$ -formula φ ($\ell \geq 2$), we can construct an equisatisfiable equality-free \mathcal{AF}^ℓ -formula ψ , also—modulo trivial logical re-arrangement—in normal form. We showed in Lemma 9 that, if φ is satisfiable over some domain, then ψ is satisfiable over the same domain. And we showed in Lemma 13 that, if ψ is satisfiable over some finite domain, then φ is satisfiable over a domain larger by a factor of at most $|I| \cdot (\ell^2 + \ell + 1)^\ell$. Moreover, a simple check shows that $\|\psi\|$ is $2^{O(\|\varphi\|)}$. It follows by a routine (if slightly fiddly) induction that any equality-free, satisfiable normal-form formula of \mathcal{AF}^ℓ ($\ell \geq 2$) has a finite model

of size $\mathfrak{t}(\ell-1, O(\|\varphi\|))$. We omit a detailed proof, because, in the next section, we strengthen the above results in two ways. First, we consider normal-form $\mathcal{AF}^{\ell+1}$ -formulas φ involving equality. Using a more elaborate definition of $\psi \in \mathcal{AF}^\ell$, we prove analogues of Lemmas 9 and 13. Secondly, as a by-product, we obtain tighter size bounds for the latter, showing that, if ψ is satisfiable over some domain, then φ is satisfiable over the same domain. This simplifies the proof of the eventual bound on model-sizes, which we give in full in Theorem 1.

§5. Adding Equality. In this section, we complete the task of establishing a small model property for each of the fragments \mathcal{AF}^ℓ , for $\ell \geq 2$, using the same strategy as employed in Sec. 4 for the equality-free sub-fragments. The difference is that, when equality is present, Lemma 10 becomes invalid. This lemma was used in the proof of Lemma 13, which constructed a model of the equality-free normal-form $\mathcal{AF}^{\ell+1}$ -formula φ from a model of the equality-free \mathcal{AF}^ℓ -formula ψ : by duplicating elements in the model of ψ , we could easily avoid clashes when selecting elements to serve as witnesses for the existential requirements of φ . When equality is present, such duplication is no longer available, thus necessitating some additional combinatorial manoeuvres. Table 3 provides a guide to important notions in this section. (Do note that they are, as of yet, undefined.)

Suppose $\mathfrak{A} \models \varphi$, where φ is an normal form $\mathcal{AF}^{\ell+1}$ -formula ($\ell \geq 2$) over some signature τ . We fix φ and τ for the remainder of this section, writing the former as in (2), again repeated here for convenience:

$$\bigwedge_{i \in I} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_i \wedge \forall \mathbf{x}_{\ell+1} \beta.$$

We employ the letters ℓ , I , β and γ_i with these denotations, and we assume without loss of generality that I is non-empty. We proceed to define an expansion \mathfrak{A}^+ of \mathfrak{A} , and simultaneously, a normal-form \mathcal{AF}^ℓ -formula ψ over the expanded signature, such that $\mathfrak{A}^+ \models \psi$; we later show that any layered model of ψ (having height ℓ) has a τ -reduct that can be elevated to a model of φ .

We take ψ to have the form

$$\psi := \text{acl}(\varphi) \wedge \psi_0 \wedge \dots \wedge \psi_5,$$

Construction of ψ from φ in the presence of =

$G_{\bar{a}}$	directed graph $(A, D_{\bar{a}} \cup E_{\bar{a}})$ of witnesses in \mathfrak{A} around \bar{a}
σ	a star; a function mapping witness $(\ell+1)$ -types to colours C
$\text{atp}(\sigma)$	the underlying ℓ -type of σ
St	the set of all stars $\sigma : \text{Atp}_{\ell+1}^\tau \leftrightarrow C$
$\text{tl}(\xi)$	the ℓ -type η s.t. $\xi \models \eta^+$
$s_\sigma(\mathbf{x}_\ell)$	atom implying \mathbf{x}_ℓ realises σ
$q_{\sigma,c}(\mathbf{x}_{\ell-1})$	atom implying there is some x s.t. $x\mathbf{x}_{\ell-1}$ realise σ and colour c
$r_{\sigma,c,\xi}(\mathbf{x}_\ell)$	atom implying x_ℓ is the ξ -witness for $x\mathbf{x}_{\ell-1}$ realising σ and c

TABLE 3. Quick reference guide for Sec 5.

where $\text{acl}(\varphi)$ is the adjacent closure of φ (featured in Lemma 7), and ψ_0, \dots, ψ_5 are \mathcal{AF}^ℓ -formulas over an expanded signature, defined below. We proceed to consider the conjuncts ψ_0, \dots, ψ_5 in turn.

The conjunct ψ_0 . The initial step in this process is rather elaborate, and has no analogue in Sec. 4. For any element $a \in A$ and any tuple $\bar{a} \in A^{\ell-1}$, let $B_{a\bar{a}}$ be a minimal set such that, for each $i \in I$, there is some $b \in B_{a\bar{a}}$ with $\mathfrak{A} \models \gamma_i[a\bar{a}b]$. Since $\mathfrak{A} \models \varphi$, such a set exists, and moreover $|B_{a\bar{a}}| \leq |I|$. We call the elements of $B_{a\bar{a}}$ the *witnesses with respect to $a\bar{a}$* . By assumption of minimality of $B_{a\bar{a}}$, there are no two distinct elements $b, b' \in B_{a\bar{a}}$ such that $\text{atp}_{\ell+1}^{\mathfrak{A}}(a\bar{a}b) = \text{atp}_{\ell+1}^{\mathfrak{A}}(a\bar{a}b')$. Thus, given an adjacent $(\ell+1)$ -type ξ such that $\text{atp}_{\ell+1}^{\mathfrak{A}}(a\bar{a}b) = \xi$ for some $b \in B_{a\bar{a}}$, we will call this b the ξ -*witness with respect to $a\bar{a}$* . (Of course, this notion depends on our particular choice of the set $B_{a\bar{a}}$.)

Having chosen the witnesses with respect to the various ℓ -tuples over A , consider now any $(\ell-1)$ -tuple \bar{a} over A . Defining the sets of ordered pairs over A

$$\begin{aligned} D_{\bar{a}} &:= \{(a, b) \mid a \neq b \text{ and } b \in B_{a\bar{a}}\} \\ E_{\bar{a}} &:= \{(a, a') \mid a \neq a' \text{ and there is some } b \in B_{a\bar{a}} \text{ s.t. } a' \in B_{b\bar{a}}\} \end{aligned}$$

we let $G_{\bar{a}}$ be the directed graph with vertices A and edges $D_{\bar{a}} \cup E_{\bar{a}}$. Thus, in this directed graph, there is an edge from a to every witness b with respect to $a\bar{a}$ (except a itself), and an edge from any element of a to any other element a' if there exists a witness b with respect to $a\bar{a}$ such that a' is a witness with respect to $b\bar{a}$. Since no ℓ -tuple has more than I witnesses, the out-degree of any vertex in $G_{\bar{a}}$ is at most $|I|^2 + |I|$. Recall that a k -*colouring* of a directed graph $G_{\bar{a}} = (V, E)$ is a function $f: V \rightarrow [0, k-1]$ satisfying $f(u) \neq f(v)$ for all $(u, v) \in E$. It is well-known that any directed graph with maximum out-degree d has a $(2d+1)$ -colouring (see e.g. [29, p. 612]). Thus, we may colour the directed graph $G_{\bar{a}}$ with colours from a set C of cardinality $2(|I|^2 + |I|) + 1$. For every $(\ell-1)$ -tuple \bar{a} , then, let some such colouring of $G_{\bar{a}}$ be fixed.

The importance of this colouring will become clearer as the proof unfolds. For the present, however, we note the following: (i) if b is one of the witnesses with respect to $a\bar{a}$, and is distinct from a , then a and b are differently coloured in $G_{\bar{a}}$; (ii) if, in addition, a' is one of the witnesses with respect to $b\bar{a}$, and is distinct from a , then a and a' are differently coloured in $G_{\bar{a}}$. It does not follow that the various witnesses with respect to $a\bar{a}$ will be differently coloured from each other in $G_{\bar{a}}$. Notice also that, for distinct $(\ell-1)$ -tuples \bar{a} and \bar{a}' , an element a might be coloured differently in the graphs $G_{\bar{a}}$ and $G_{\bar{a}'}$; this is true even if $\bar{a}' = \bar{a}$.

Now let us treat the elements of C as ℓ -ary predicates, which we interpret in our expansion \mathfrak{A}^+ of \mathfrak{A} . Specifically, for all $a \in A$, $\bar{a} \in A^{\ell-1}$ and $c \in C$, we declare that $\mathfrak{A}^+ \models c[a\bar{a}]$ just in case a is assigned the colour c in the colouring of $G_{\bar{a}}$, and we write $\text{col}^{\mathfrak{A}^+}[a\bar{a}]$ to denote c . Note that, defining

$$\psi_0 := \forall \mathbf{x}_\ell \left(\bigvee_{c \in C} c(\mathbf{x}_\ell) \wedge \bigwedge_{\substack{c \neq c' \\ c, c' \in C}} \left(\neg c(\mathbf{x}_\ell) \vee \neg c'(\mathbf{x}_\ell) \right) \right),$$

we have $\mathfrak{A}^+ \models \psi_0$. Conversely, if \mathfrak{B} is any model of ψ_0 , every ℓ -tuple over B is assigned a unique colour from C in the obvious way.

The conjunct ψ_1 . Here we can simply repeat material from Sec. 4. For each s ($2 < 2s + 1 \leq \ell$) we introduce a fresh $(2s+1)$ -ary predicate d_{2s+1} , and declare that a $(2s+1)$ -tuple over A satisfies d_{2s+1} in the expansion \mathfrak{A}^+ just in case it is a palindrome. In addition, we define the \mathcal{AF}^ℓ -formula

$$\psi_1 := \bigwedge_{2 < 2s+1 \leq \ell} \forall \mathbf{x}_{s+1} d_{2s+1}(x_1 \cdots x_s x_{s+1} x_s \cdots x_1).$$

Thus, $\mathfrak{A}^+ \models \psi_1$, and, conversely, if \mathfrak{B} is any structure such that $\mathfrak{B} \models \psi_1$, and $\bar{c} \in B^{2s+1}$ is a palindrome ($2 < 2s + 1 \leq \ell$), then $\mathfrak{B} \models d_{2s+1}[\bar{c}]$. As before, we write $\delta_D := \bigwedge \{d_{j-i+1}(x_i \cdots x_j) \mid \langle i, j \rangle \in D\}$ where D is any set of pairs of integers $\langle i, j \rangle$ for $1 \leq i < j \leq \ell$ with $j-i+1$ odd and greater than 2.

The conjunct ψ_2 . The treatment of this conjunct is more elaborate than in Sec. 4, and involves the colouring of ℓ -tuples encountered above. However, the essential function of securing witnesses and imposing universal constraints is the same. We remind the reader that, if \bar{a} is a tuple, then \bar{a} denotes its reversal. Observe that each ℓ -tuple $a\bar{a}$ over A gives rise to a partial function $\sigma: \text{Atp}_{\ell+1}^\tau \hookrightarrow C$ mapping adjacent $(\ell+1)$ -types over τ to colours, namely,

$$\sigma(\xi) = \begin{cases} \text{col}^{\mathfrak{A}^+}[b\bar{a}] & \text{if } b \text{ is the (unique) } \xi\text{-witness in } \mathfrak{A} \text{ with respect to } a\bar{a}, \\ \text{undefined} & \text{if } a\bar{a} \text{ does not have a } \xi\text{-witness in } \mathfrak{A}. \end{cases}$$

Thus, for every adjacent $(\ell+1)$ -type ξ for which $a\bar{a}$ has a ξ -witness b in \mathfrak{A} , the value $\sigma(\xi)$ tells us the colour of the ‘reversed’ ℓ -tuple $b\bar{a}$ in \mathfrak{A}^+ . We denote the domain of σ (i.e. the set of $\xi \in \text{Atp}_{\ell+1}^\tau$ for which $\sigma(\xi)$ is defined) by $\text{dom}(\sigma)$. Since I is non-empty, the ℓ -tuple $a\bar{a}$ has at least one witness, whence $\text{dom}(\sigma)$ is also non-empty. We call σ the *star of $a\bar{a}$* in \mathfrak{A}^+ , and denote it $\text{st}^{\mathfrak{A}^+}[a\bar{a}]$. Always remember that the elements of $\text{Atp}_{\ell+1}^\tau$ are adjacent types over the *original* signature τ of \mathfrak{A} , not over any expanded signature.

A simple check with reference to the formula φ given above verifies that σ satisfies the following conditions:

1. there exists a (unique) adjacent ℓ -type ζ over τ such that, for all $\xi \in \text{dom}(\sigma)$, we have $\xi \models \zeta$.
2. for every $\xi \in \text{dom}(\sigma)$ there exists $i \in I$ such that $\xi \models \gamma_i$;
3. for every $i \in I$ there is exactly one $\xi \in \text{dom}(\sigma)$ such that $\xi \models \gamma_i$; and
4. for every $\xi \in \text{dom}(\sigma)$, $\xi \models \hat{\beta}$ (remember that $\hat{\beta} := \beta \wedge \beta^{-1}$).

Condition 1 is verified by setting $\zeta = \text{atp}^{\mathfrak{A}^+}[a\bar{a}]$; it is obvious that this is the unique adjacent ℓ -type with the required properties; we call ζ the *underlying adjacent type* of σ and denote it $\text{atp}(\sigma)$. The remaining conditions are immediate from the properties of witnesses. Accordingly, we call any partial function $\sigma: \text{Atp}_{\ell+1}^\tau \hookrightarrow C$ satisfying conditions 1–4 above a *star*, and we denote the set of all such stars as St . Thus, for any $(\ell+1)$ -tuple over A , we have $\text{st}^{\mathfrak{A}^+}[a\bar{a}] \in \text{St}$. We remark that the notion of a star depends on the formula φ (and on the parameters $\ell, C, \beta, I, \gamma_i$ and τ associated with φ). Since, however, φ may be considered fixed for the present, we suppress these parameters to avoid notational clutter.

Now we are ready to fix some additional predicates in the expansion \mathfrak{A}^+ . Since we have settled the interpretations of the predicates C in \mathfrak{A}^+ , the adjacent star-type $\text{st}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ of any ℓ -tuple $a\bar{a}$ is unaffected by these additional predicates: in particular, the domain of σ consists of adjacent $(\ell+1)$ -types ξ over the *original signature* τ ; and the values $\sigma(\xi)$ are determined by the interpretations of the predicates C in \mathfrak{A}^+ . With this in mind, for every $\zeta \in \text{Atp}_\ell^\tau$, we introduce the $(\ell-1)$ -ary predicate p_ζ familiar from Sec.4, declaring $\mathfrak{A}^+ \models p_\zeta[\bar{a}]$ just in case, for some $a \in A$, $\text{atp}^{\mathfrak{A}^+}[a\bar{a}] = \zeta$. Thus, $\mathfrak{A}^+ \models \psi_{2,0}$, where

$$\psi_{2,0} := \bigwedge_{\zeta \in \text{Atp}_\ell^\tau} \forall \mathbf{x}_\ell (\zeta \rightarrow p_\zeta(x_2 \cdots x_\ell)).$$

In addition, for every $\sigma \in \text{St}$, we introduce a new ℓ -ary predicate s_σ , and declare $\mathfrak{A}^+ \models s_\sigma[a\bar{a}]$ if and only if $\text{st}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$, for any ℓ -tuple $a\bar{a}$ over A . It is then easy to verify that $\mathfrak{A}^+ \models \psi_{2,1} \wedge \psi_{2,2}$, where

$$\begin{aligned} \psi_{2,1} &:= \forall \mathbf{x}_\ell \bigvee_{\sigma \in \text{St}} s_\sigma(\mathbf{x}_\ell) \wedge \forall \mathbf{x}_\ell \bigwedge_{\substack{\sigma \neq \sigma' \\ \sigma, \sigma' \in \text{St}}} (\neg s_\sigma(\mathbf{x}_\ell) \vee \neg s_{\sigma'}(\mathbf{x}_\ell)) \\ \psi_{2,2} &:= \forall \mathbf{x}_\ell \bigvee_{\sigma \in \text{St}} (s_\sigma(\mathbf{x}_\ell) \rightarrow \text{atp}(\sigma)). \end{aligned}$$

Indeed, the first of these formulas states that any ℓ -tuple satisfies exactly one of the predicates s_σ , and the second, that its adjacent type in \mathfrak{A} is given by $\text{atp}(\sigma)$. (Recall in this connection that $\text{atp}(\sigma)$ is an ℓ -type, and thus a formula with variables \mathbf{x}_ℓ .) Further, for every $\sigma \in \text{St}$ and every $c \in C$, we introduce a new $(\ell-1)$ -ary predicate $q_{\sigma,c}$, and declare $\mathfrak{A}^+ \models q_{\sigma,c}[\bar{a}]$ just in case there is some $a \in A$ such that $\text{st}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ and $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$. Thus, $q_{\sigma,c}$ identifies tails of ℓ -tuples whose star in \mathfrak{A}^+ is σ and whose colour in \mathfrak{A}^+ is c . It is thus immediate that $\mathfrak{A}^+ \models \psi_{2,3}$, where

$$\psi_{2,3} := \forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{St}} \bigwedge_{c \in C} \left((s_\sigma(\mathbf{x}_\ell) \wedge c(\mathbf{x}_\ell)) \rightarrow q_{\sigma,c}(x_2 \cdots x_\ell) \right).$$

Still proceeding with the construction of ψ_2 , for every $\sigma \in \text{St}$, every $c \in C$, and every $\xi \in \text{dom}(\sigma)$, we introduce a new ℓ -ary predicate $r_{\sigma,c,\xi}$, and fix its interpretation in \mathfrak{A}^+ as follows. Take any $(\ell-1)$ -tuple \bar{a} over A . If, on the one hand, $\mathfrak{A}^+ \models q_{\sigma,c}[\bar{a}]$, then select some $a \in A$ such that $\text{st}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ and $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$. By the interpretation of $q_{\sigma,c}$ in \mathfrak{A}^+ , this is possible. Now, for each $\xi \in \text{dom}(\sigma)$, let b be the ξ -witness for $a\bar{a}$, and set $\mathfrak{A}^+ \models r_{\sigma,c,\xi}[\bar{a}b]$. If, on the other hand, $\mathfrak{A}^+ \not\models q_{\sigma,c}[\bar{a}]$, then do nothing in respect of the tuple \bar{a} . By carrying out this procedure for every $(\ell-1)$ -tuple \bar{a} over A , we thus fix the extensions of the predicates $r_{\sigma,c,\xi}$ in \mathfrak{A}^+ . Informally, it helps to read the atom $r_{\sigma,c,\xi}(\mathbf{x}_\ell)$ as “ x_ℓ wants to be the ξ -witness with respect to the tuple $x'\mathbf{x}_{\ell-1}$ for a particular x' such that $x'\mathbf{x}_{\ell-1}$ has star-type σ and colour c ”. Since $\xi \in \text{dom}(\sigma)$ implies that a ξ -witness exists, we have $\mathfrak{A}^+ \models \psi_{2,4}$, where

$$\psi_{2,4} := \forall \mathbf{x}_{\ell-1} \exists x_\ell \bigwedge_{\sigma \in \text{St}} \bigwedge_{c \in C} \bigwedge_{\xi \in \text{dom}(\sigma)} (q_{\sigma,c}(\mathbf{x}_{\ell-1}) \rightarrow r_{\sigma,c,\xi}(\mathbf{x}_\ell)).$$

Let us return to our ℓ -tuple $a\bar{a}$ with $\text{st}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$, $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$ and ξ -witness b . Let $\text{tl}(\xi)$ denote the unique adjacent ℓ -type η over τ such that $\xi \models \eta^+$. Thus, $\text{atp}^{\mathfrak{A}}[\bar{a}b] = \text{tl}(\xi)$. Moreover, from the fact that $\text{st}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ and b is the ξ -witness with respect to $a\bar{a}$, we have $\text{col}^{\mathfrak{A}^+}[b\bar{a}] = \sigma(\xi)$. Thus, $\mathfrak{A}^+ \models \psi_{2,5}$, where

$$\psi_{2,5} := \forall \mathbf{x}_{\ell-1} \bigwedge_{\sigma \in \text{St}} \bigwedge_{c \in C} \bigwedge_{\xi \in \text{dom}(\sigma)} (r_{\sigma,c,\xi}(\mathbf{x}_{\ell}) \rightarrow \text{tl}(\xi) \wedge (\sigma(\xi))(x_{\ell} \cdots x_1)).$$

Recall in this regard that $\text{tl}(\xi)$ is an adjacent ℓ -type (hence a formula with free variables \mathbf{x}_{ℓ}), and $\sigma(\xi) \in C$ is an ℓ -ary predicate of τ .

More can be said considering the ℓ -tuple $\bar{a}b$ and its reversal, $b\bar{a}$. This latter tuple has some star in \mathfrak{A}^+ , say $\text{st}^{\mathfrak{A}^+}[b\bar{a}] = \sigma'$. Consider, then, any adjacent $(\ell+1)$ -type $\xi' \in \text{dom}(\sigma')$. Thus, $b\bar{a}$ has a ξ' -witness in \mathfrak{A} , say b' . If $b' = a$, then $\text{atp}^{\mathfrak{A}}[a\bar{a}b]$ and $\text{atp}^{\mathfrak{A}}[b\bar{a}b']$ are mutually inverse $(\ell+1)$ -types, and we have $\xi' = \xi^{-1}$. If, on the other hand, $b' \neq a$, then $\langle a, b' \rangle$ is an edge of the directed graph $G_{\bar{a}}$, whence $\text{col}^{\mathfrak{A}^+}[a\bar{a}] \neq \text{col}^{\mathfrak{A}^+}[b'\bar{a}]$, by construction of \mathfrak{A}^+ . Thus, $\mathfrak{A}^+ \models \psi_{2,6}$, where

$$\psi_{2,6} := \forall \mathbf{x}_{\ell} \bigwedge_{\sigma \in \text{St}} \bigwedge_{c \in C} \bigwedge_{\xi \in \text{dom}(\sigma)} \left(r_{\sigma,c,\xi}(\mathbf{x}_{\ell}) \rightarrow \bigvee \{ s_{\sigma'}(x_{\ell} \cdots x_1) \mid \sigma' \in \text{St} \text{ and, for all } \xi' \in \text{dom}(\sigma'), \xi' \neq \xi^{-1} \Rightarrow \sigma'(\xi') \neq c \} \right).$$

That is to say: if b is the ξ -witness with respect to $a\bar{a}$ and b' the ξ' -witness with respect to $b\bar{a}$, then either ξ and ξ' are mutually inverse types in \mathfrak{A} or else the tuples $a\bar{a}$ and $b'\bar{a}$ are differently coloured in \mathfrak{A}^+ .

We complete the construction of ψ_2 with a somewhat simpler constraint concerning the predicates $r_{\sigma,c,\xi}$. Consider again any $(\ell-1)$ -tuple \bar{a} over A , and suppose that, for some $\sigma \in \text{St}$, some $c \in C$ and some pair of distinct $(\ell+1)$ -types $\xi, \xi' \in \text{dom}(\sigma)$, there exist elements b and b' such that $\mathfrak{A}^+ \models r_{\sigma,c,\xi}[\bar{a}b]$ and $\mathfrak{A}^+ \models r_{\sigma,c,\xi'}[\bar{a}b']$. From the construction of \mathfrak{A}^+ , b and b' must have been chosen as ξ - and ξ' -witnesses, respectively, with respect to a tuple $a\bar{a}$ for some particular element a . It follows that b and b' must be distinct. That is, $\mathfrak{A}^+ \models \psi_{2,7}$, where

$$\psi_{2,7} := \forall \mathbf{x}_{\ell} \bigwedge_{\sigma \in \text{St}} \bigwedge_{c \in C} \bigwedge_{\xi, \xi' \in \text{dom}(\sigma)}^{\xi \neq \xi'} (r_{\sigma,c,\xi}(\mathbf{x}_{\ell}) \rightarrow \neg r_{\sigma,c,\xi'}(\mathbf{x}_{\ell})).$$

Conversely, if \mathfrak{B} is any structure making $\psi_{2,7}$ true, and \bar{a} an $(\ell-1)$ -tuple over B , then we cannot have the same element b such that $\mathfrak{B} \models r_{\sigma,c,\xi}[\bar{a}b]$ and $\mathfrak{B} \models r_{\sigma,c,\xi'}[\bar{a}b]$ for different ξ and ξ' in the domain of σ .

Setting $\psi_2 := \psi_{2,0} \wedge \cdots \wedge \psi_{2,7}$, we have thus established that $\mathfrak{A}^+ \models \psi_2$.

The conjuncts ψ_3 and ψ_4 . Here, we can again recapitulate the ideas of Sec. 4, though in a slightly different guise. As before, we take \mathbf{D}_k° to denote the set of all pairs $\langle i, j \rangle$ for $1 \leq i < j \leq k$ such that $j-i+1$ is greater than 2 and odd. Fix some subset $D \subseteq \mathbf{D}_{\ell-1}^{\circ}$, and suppose \bar{a} is an $(\ell-1)$ -tuple over A such that $\mathfrak{A}^+ \models \delta_D[\bar{a}]$. By construction of \mathfrak{A}^+ , the defect set of \bar{a} includes D , whence, for any elements $a, b \in A$, the defect set of $a\bar{a}b$ certainly includes D^+ . It follows that, if $\sigma \in \text{St}$, $c \in C$ and $\xi \in \text{dom}(\sigma)$, but with $\partial\xi$ not D^+ -compatible, then there cannot exist $b \in A$ such that $\mathfrak{A} \models r_{\sigma,c,\xi}[\bar{a}b]$. For otherwise, by construction

of \mathfrak{A}^+ , there exists a such that $\text{st}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$, and b is the ξ -witness with respect to $a\bar{a}$, whence $\text{atp}^{\mathfrak{A}^+}[a\bar{a}b] = \xi$, contradicting the first statement of Lemma 8. We have thus proved $\mathfrak{A}^+ \models \psi_3$, where

$$\psi_3 := \bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^0} \forall \mathbf{x}_\ell (\delta_D \rightarrow \bigwedge \{ \neg r_{\sigma, c, \xi}(\mathbf{x}_\ell) \mid \sigma \in \text{St}, c \in C, \xi \in \text{dom}(\sigma), \partial\xi \text{ not } D^+ \text{-compatible} \}).$$

For ψ_4 , we recapitulate the formula from Sec. 4, namely

$$\psi_4 := \bigwedge_{\zeta \in \text{Atp}_\ell^\tau} \bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^0} \forall \mathbf{x}_\ell \left((\delta_D \wedge p_\zeta(\mathbf{x}_{\ell-1})) \rightarrow \bigvee \{ \eta \in \text{Atp}_\ell^\tau \mid (\zeta \wedge \hat{\beta} \wedge \eta^+) \text{ is } D^+ \text{-consistent} \} \right).$$

And by the same reasoning as in Sec. 4, we have $\mathfrak{A}^+ \models \psi_4$.

The conjunct ψ_5 . Our final conjunct again has no analogue in Sec. 4, and concerns extra conditions we need to impose on certain *non-primitive* $(\ell+1)$ -tuples in models of ψ . We need to ensure that, when reconstructing a model of φ from such structures, we do not attempt to assign these tuples incompatible adjacent types. We require the following simple lemma regarding words. Recall that $\bar{\mathbf{A}}_k^\ell$ denotes the set of adjacent functions $f: [1, \ell] \rightarrow [1, k]$ such that $f(\ell) = k$. We say that an ℓ -tuple \bar{b} is *terminal* if it can be written $\bar{b} = \bar{d}^f$ for some k -tuple \bar{d} with $k < \ell$, and some $f \in \bar{\mathbf{A}}_k^\ell$. As explained in Sec. 2, for any $f \in \bar{\mathbf{A}}_k^\ell$, a leg of f is a maximal interval $[i, j] \subseteq [1, \ell]$ such that $f(h+1) - f(h)$ is constant for all h ($i \leq h < j$); such an interval corresponds to a straight-line segment of the graph of f .

LEMMA 14. *Let \bar{b} be an ℓ -tuple over some set and b an element of that set. Then at least one of the following holds: (i) $\bar{b}b$ is primitive; (ii) b is the last element of \bar{b} ; (iii) $\bar{b}b$ has a suffix which is an odd-length, non-trivial palindrome; (iv) \bar{b} is terminal.*

PROOF. Suppose that $\bar{b}b$ is not primitive. If \bar{b} has an immediately repeated letter, then it is certainly terminal: indeed $\bar{b} = \bar{a}c\bar{c}\bar{d}$ is generated from $\bar{a}\bar{c}\bar{d}$ via a function $f \in \bar{\mathbf{A}}_{\ell-1}^\ell$ which pauses for one step on the letter c . Hence we may assume that there are no immediately repeated letters in \bar{b} . Furthermore, if b is not the last element of \bar{b} , then the whole of $\bar{b}b$ has no immediately repeated letters. Since $\bar{b}b$ is not primitive, we have $\bar{b}'b = \bar{c}^f$ for some k -tuple \bar{b}' and some $f \in \bar{\mathbf{A}}_k^m$ with at least two legs (i.e. maximal strictly ascending or descending intervals). If the final leg of f is shortest, then some suffix of $\bar{b}b$ is a non-trivial palindrome, and this palindrome must have odd length, since there are no immediately repeated letters. Otherwise $\bar{b}b$ has either of the forms $\bar{c}\bar{d}\bar{d}\bar{c}\bar{e}\bar{b}$ or $\bar{a}\bar{c}\bar{d}\bar{d}\bar{c}\bar{d}\bar{e}\bar{b}$, depending on whether the shortest leg is initial or internal. In the former case, $\bar{b} = (\bar{d}\bar{c}\bar{e})^g$ for some final adjacent function g ; in the latter, $\bar{b} = (\bar{a}\bar{c}\bar{d}\bar{d}\bar{e})^h$ for some final adjacent function h . In both cases, \bar{b} is terminal. \dashv

We remark that, in the case were $\bar{b}b$ has a suffix which is an odd-length non-trivial palindrome, that palindrome may be the whole of $\bar{b}b$. The cases of Lemma 14 correspond to conditions on the adjacent type of the $(\ell+1)$ -tuple in

question. Say that an adjacent $(\ell+1)$ -type ξ is *palindromic* if, for any $\mathcal{AF}^{[\ell+1]}$ -atom α (over the relevant signature), $\xi \models \alpha$ implies $\xi \models \alpha(x_{\ell+1} \cdots x_1)$. Evidently, if an $(\ell+1)$ -tuple is a palindrome, then its adjacent type in any structure is palindromic. Say that ξ is *blunt* if, for any $\mathcal{AF}^{[\ell+1]}$ -atom α , $\xi \models \alpha$ implies $\xi \models \alpha(\mathbf{x}_\ell x_\ell)$. If an $(\ell+1)$ -tuple has the same last two elements, then its adjacent type in any structure is blunt. Say that ξ is *s-hooked* (for s satisfying $2 < 2s + 1 < \ell+1$) if, for any $\mathcal{AF}^{[\ell+1]}$ -atom α we have that $\xi \models \alpha$ implies $\xi \models \alpha(x_1 \cdots x_{\ell+1-s} x_{\ell-s} \cdots x_{\ell+1-2s})$. If an $(\ell+1)$ -tuple has a proper suffix that is a non-trivial palindrome of length $2s+1$, then its adjacent type in any structure is s -hooked. Observe the strict inequality $2s + 1 < \ell+1$ governing the parameter s in this last definition: if ξ is palindromic (and $\ell + 1$ is odd), we do not say that ξ is $(\ell/2)$ -hooked.

Now let \bar{a} be an $(\ell-1)$ -tuple over A and $b \in A$, and suppose $\mathfrak{A}^+ \models r_{\sigma,c,\xi}[\bar{a}b]$. By the construction of \mathfrak{A}^+ , there exists $a \in A$ such that b is the ξ -witness with respect to $a\bar{a}$, and, moreover, $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$. If ξ is palindromic, blunt or s -hooked for some s ($2 < 2s + 1 < \ell+1$), then the $(\ell+1)$ -tuple $a\bar{a}b$ exhibits certain properties, which we proceed to describe. Consider first the case where $\ell+1$ is odd and ξ is not palindromic, and suppose in addition that $\mathfrak{A}^+ \models d_{\ell-1}[\bar{a}]$. By the construction of \mathfrak{A}^+ again, \bar{a} is a palindrome, i.e. $\bar{a} = \bar{a}$. It follows that $a \neq b$, since otherwise, $a\bar{a}b$ is a palindrome, contradicting the assumption that $\xi = \text{atp}^{\mathfrak{A}^+}[a\bar{a}b]$ is not palindromic. And since b is a witness with respect to $a\bar{a}$ with $a \neq b$, the ordered pair $\langle a, b \rangle$ is an edge of the directed graph $G_{\bar{a}} = G_{\bar{a}}$, so that by the construction of \mathfrak{A}^+ , we have $c = \text{col}^{\mathfrak{A}^+}[a\bar{a}] \neq \text{col}^{\mathfrak{A}^+}[b\bar{a}] = \text{col}^{\mathfrak{A}^+}[b\bar{a}]$. Thus we have shown that $\mathfrak{A}^+ \models \psi_{5,1}$, where, for $(\ell+1)$ odd,

$$\psi_{5,1} := \forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{St}} \bigwedge_{c \in C} \bigwedge_{\substack{\xi \text{ not palindromic} \\ \xi \in \text{dom}(\lambda)}} \left(r_{\lambda,c,\xi}(\mathbf{x}_\ell) \wedge d_{\ell-1}(\mathbf{x}_{\ell-1}) \rightarrow \neg c(x_\ell \cdots x_1) \right),$$

and for $(\ell+1)$ even, $\psi_{5,1} := \top$. We remark that even-length palindromes have immediately repeated letters in the middle, which obviates—as we shall see later—the need for an analogue of the odd-length case.

Second, consider the case where ξ is not blunt. Here, we do not need to add any conjuncts to ψ , since the consistency of ξ requires that it contains the inequality literal $x_\ell \neq x_{\ell+1}$. Hence, $\text{atp}^{\mathfrak{A}^+}[\bar{a}b]$ contains the inequality literal $x_{\ell-1} \neq x_\ell$, which is all the information we shall require.

Third, consider the case where, for any s in the range $2 < 2s + 1 < \ell+1$, ξ is not s -hooked. Since $\xi = \text{atp}^{\mathfrak{A}^+}[a\bar{a}b]$, it follows that $a\bar{a}b$ has no proper suffix of length $2s + 1$ that is a palindrome, and therefore, $\bar{a}b$ has no suffix of the same length that is a palindrome. Thus we have shown that $\mathfrak{A}^+ \models \psi_{5,2}$, where

$$\psi_{5,2} := \forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{St}} \bigwedge_{c \in C} \bigwedge_{s=1}^{s \leq (\ell-1)/2} \bigwedge_{\xi \in \text{dom}(\lambda)} \left(r_{\lambda,c,\xi}(\mathbf{x}_\ell) \rightarrow \neg d_{2s+1}(x_{\ell-2s} \cdots x_\ell) \right).$$

Writing $\psi_5 := \psi_{5,1} \wedge \psi_{5,2} \wedge \psi_{5,2}$, we have shown that $\mathfrak{A}^+ \models \psi_5$. This completes the definition of the formula ψ .

Summarizing the above discussion, and recalling that, by Lemma 7, $\varphi \models \text{acl}(\varphi)$, we have:

LEMMA 15. *Suppose $\mathfrak{A} \models \varphi$. Then we can expand \mathfrak{A} to a model $\mathfrak{A}^+ \models \psi$.*

Having defined ψ and established Lemma 15, we establish a converse in the form of the following lemma.

LEMMA 16. *Suppose $\mathfrak{A} \models \psi$. Then we can construct a model $\mathfrak{B} \models \varphi$ over the same domain.*

PROOF. Since $\psi \in \mathcal{AF}^\ell$, by Lemma 5 we may take \mathfrak{A} to be a layered structure of the height ℓ . Setting \mathfrak{A}^- to be the contraction of \mathfrak{A} to the signature τ (the original signature of φ), we proceed to elevate \mathfrak{A}^- to a model $\mathfrak{B} \models \varphi$ (with the same domain, A) by setting the interpretations of the predicates in τ with respect to the primitive $(\ell+1)$ -tuples over A .

Let a be any element of A and \bar{a} any $(\ell-1)$ -tuple over A . By ψ_0 , $a\bar{a}$ has a unique colour, say $c = \text{col}^{\mathfrak{A}}[a\bar{a}]$, and by $\psi_{2,1}$, there is a unique $\sigma \in \text{St}$ such that $\mathfrak{A} \models s_\sigma[a\bar{a}]$. Let $\zeta = \text{atp}(\sigma)$ be the underlying ℓ -type of σ , and let us fix, for the moment, any $\xi \in \text{dom}(\sigma)$. Thus, $\zeta = \xi|_{[1,\ell]}$. By $\psi_{2,2}$, $\text{atp}^{\mathfrak{A}^-}[a\bar{a}] = \zeta$. Moreover, by $\psi_{2,3}$, $\mathfrak{A} \models q_{\sigma,c}[\bar{a}]$, and hence, by the fact that $\mathfrak{A} \models \psi_{2,4}$, there exists $b \in A$ such that $\mathfrak{A} \models r_{\sigma,c,\xi}[\bar{a}b]$. Let $\eta = \text{tl}(\xi)$; it follows from $\psi_{2,5}$ that $\text{atp}^{\mathfrak{A}^-}[\bar{a}b] = \eta$ and $\text{col}^{\mathfrak{A}}[\bar{a}b] = \sigma(\xi)$. The intention is that we should set the interpretations of the predicates in the structure \mathfrak{B} in such a way that the $(\ell+1)$ -tuple $a\bar{a}b$ is assigned the adjacent type ξ . The various other conjuncts of ψ will ensure that this can be done consistently.

Suppose on the one hand that the $(\ell+1)$ -tuple $a\bar{a}b$ is primitive. Thus, no prefix or suffix of $a\bar{a}b$ is a non-trivial palindrome, and $a\bar{a}b$ contains no immediately repeated letters. Hence, any defect $\langle i, j \rangle$ of $a\bar{a}b$ satisfies $2 \leq i < j \leq \ell$, with $j-i+1$ greater than 2 and odd. Letting D now be the set of defects $\langle i, j \rangle$ of the $(\ell-1)$ -tuple \bar{a} (so that $D \subseteq \mathbf{D}_{\ell-1}^\circ$), we see that the set of defects of $a\bar{a}b$ is given by $D^+ = \{\langle i+1, j+1 \rangle \mid \langle i, j \rangle \in D\}$. By ψ_1 , we have that $\mathfrak{A} \models \delta_D[\bar{a}]$; and by ψ_3 , bearing in mind that $\mathfrak{A} \models r_{\sigma,c,\xi}[\bar{a}b]$, we have that $\partial\xi$ is D^+ -compatible. Lemma 8 thus ensures that it is meaningful to set $\text{itp}_{\ell+1}^{\mathfrak{B}}(a\bar{a}b) = \partial\xi$. Moreover, since $\text{atp}^{\mathfrak{A}^-}[a\bar{a}] = \zeta$, $\text{atp}^{\mathfrak{A}^-}[\bar{a}b] = \eta$, and $\xi = \zeta \cup \eta^+ \cup \partial\xi$, we have $\text{atp}_{\ell+1}^{\mathfrak{B}}(a\bar{a}b) = \xi$. We remark that the adjacent type of the primitive $(\ell+1)$ -tuple $b\bar{a}a$ is also set in this process, but no other primitive $(\ell+1)$ -tuples have their adjacent types set. Observe that, since $\xi \models \hat{\beta}$, the newly-set adjacent types of $a\bar{a}b$ and $b\bar{a}a$ will not violate β .

Suppose on the other hand that that $a\bar{a}b$ is not primitive. We show that, in that case, \mathfrak{A} already provides a ξ -witnesses with respect to $a\bar{a}$, so that there is nothing to do. Here, we make use of the fact that $\mathfrak{A} \models \text{acl}(\varphi) \wedge \psi_5$. By Lemma 14, either b is the last element of $a\bar{a}$, or $a\bar{a}b$ has a suffix that is a non-trivial odd-length palindrome, or $a\bar{a}$ is terminal. The case where $a\bar{a}$ is terminal is easily dealt with. There exists a k -tuple \bar{d} and $f \in \bar{\mathbf{A}}_k^\ell$ such that $a\bar{a} = \bar{d}^f$ for some k ($2 \leq k < \ell$). Defining, as before, $f^+ = f \cup \langle k+1, \ell+1 \rangle$, by $\text{acl}(\varphi)$ there exists, for each $i \in I$, some $b' \in A$ such that $\mathfrak{A} \models \gamma_i[(\bar{d}b')^{f^+}]$, i.e. $\mathfrak{A} \models \gamma_i[a\bar{a}b']$. In effect,

we are discarding our originally chosen element b , since the required witness is already present.

Consider next the case where b is the last element of \bar{a} . We have already argued that $\text{atp}^{\mathfrak{A}^-}[\bar{a}b] = \eta = \text{tl}(\xi)$, whence $\text{tl}(\xi) \models x_{\ell-1} = x_\ell$, whence $\xi \models x_\ell = x_{\ell+1}$. Thus, ξ is blunt, by consistency of ξ . But in that case, writing $\bar{a} = a_1 \cdots a_{\ell-1}$, we have $\mathfrak{A}^- \models \zeta[aa_1 \cdots a_{\ell-1}]$ implies $\mathfrak{A}^- \models \xi[aa_1 \cdots a_{\ell-1}a_{\ell-1}]$, that is, $\mathfrak{A}^- \models \xi[a\bar{a}b]$. Thus, our chosen element b is already a ξ -witness with respect to $a\bar{a}$, without our having to do anything.

Consider finally the case where $a\bar{a}b$ is not terminal and has a suffix that is an odd-length, non-trivial palindrome. Suppose, on the one hand, that the suffix in question is the *whole* of $a\bar{a}b$ —that is, $\ell+1$ is odd, and $a\bar{a}b$ has the form $aa_1 \cdots a_{\ell/2}a_{\ell/2-1} \cdots a_1a$. Thus, $a = b$ and $\bar{a} = \tilde{a}$. Hence, $a\bar{a} = b\tilde{a}$, whence $\text{col}^{\mathfrak{A}^-}[b\tilde{a}] = \text{col}^{\mathfrak{A}^-}[a\tilde{a}] = c$. It follows by ψ_1 that $\mathfrak{A} \models d_{\ell-1}[\tilde{a}]$, and thence, by $\psi_{5,1}$ that ξ is palindromic. Thus, $\text{atp}^{\mathfrak{A}^-}(aa_1 \cdots a_{\ell/2}) = \zeta \upharpoonright_{[1, \ell/2+1]} = \xi \upharpoonright_{[1, \ell/2+1]}$, whence $\text{atp}^{\mathfrak{A}^-}(aa_1 \cdots a_{\ell/2}a_{\ell/2-1} \cdots a_1a) = \xi$. Again, our chosen element $b = a$ is already a ξ -witness with respect to $a\bar{a}$.

On the other hand, suppose that $a\bar{a}b$ has a *proper* suffix that is a non-trivial, odd-length palindrome. Thus, $a\bar{a}b$ is of the form $aa_1 \cdots a_{\ell-s}a_{\ell-(s+1)} \cdots a_{\ell-2s}$ (for some s with $2 < 2s+1 \leq \ell$). But by ψ_1 , $\mathfrak{A} \models d_s[a_{\ell-2s} \cdots a_{\ell-s}a_{\ell-(s+1)} \cdots a_{\ell-2s}]$. Hence, by $\psi_{5,2}$, we have that ξ is s -hooked. As in the previous case, it follows that $b = a_{(\ell-2s)}$ is already a ξ -witness with respect to $a\bar{a}$. This completes the process of finding a ξ -witness with respect to $a\bar{a}$.

Keeping a and \bar{a} fixed for the moment, carry out the above assignments for all $\xi \in \text{dom}(\sigma)$, choosing, for each such ξ , a ξ -witness b_ξ with respect to $a\bar{a}$. It follows from $\psi_{2,7}$ that the various elements b_ξ must be distinct, so no clashes can arise. This completes the process of finding all the required witnesses with respect to $a\bar{a}$. From the properties of the elements of St , we see that, however \mathfrak{B} is completed, the ℓ -tuple $a\bar{a}$ will satisfy $\exists x_{\ell+1} \gamma_i$ in \mathfrak{B} for every $i \in I$. Moreover, each of the newly-fixed primitive $(\ell+1)$ -tuples (and their reversals) satisfies β .

Now vary a and \bar{a} freely. We must show that, for a' and \bar{a}' with $a\bar{a} \neq a'\bar{a}'$, the chosen witnesses b and b' can never lead to a clash. Suppose then that some m -tuple \bar{d} is assigned (or not) to the extension of an m -ary predicate p when defining the adjacent $(\ell+1)$ -types of both of the primitive tuples $a\bar{a}b$ and $a'\bar{a}'b'$. Then \bar{d} is generated by both $a\bar{a}b$ and $a'\bar{a}'b'$, and there exist adjacent $(\ell+1)$ -types ξ and ξ' such that b is selected as the ξ -witness with respect to $a\bar{a}$ and b' is selected as the ξ' -witness with respect to $a'\bar{a}'$. Since, by assumption, $a\bar{a} \neq a'\bar{a}'$, Theorem 2 implies $a\bar{a}b$ is the reversal of $a'\bar{a}'b'$, so that $a' = b$, $b' = a$ and $\bar{a}' = \tilde{a}$; henceforth, we shall write $b\tilde{a}$ in preference to $a'\bar{a}'$. It suffices to show that $\xi' = \xi^{-1}$, because in that case there can be no clash in the assignment of \bar{d} to the extension of the predicate p .

Suppose, for contradiction that $\xi' \neq \xi^{-1}$. By $\psi_{2,1}$, let σ and σ' be the unique elements of St such that $\mathfrak{A} \models s_\sigma[a\bar{a}]$ and $\mathfrak{A} \models s_{\sigma'}[b\tilde{a}]$; and by ψ_0 , let c and c' be the unique elements of C such that $\mathfrak{A} \models c[a\bar{a}]$ and $\mathfrak{A} \models c'[b\tilde{a}]$. Since b was chosen as a ξ -witness for $a\bar{a}$, we have $\mathfrak{A} \models r_{\sigma,c,\xi}[a\bar{a}]$, and similarly $\mathfrak{A} \models r_{\sigma',c',\xi'}[b\tilde{a}]$. Recalling that $\mathfrak{A} \models r_{\sigma,c,\xi}[a\bar{a}]$ and $\xi' \neq \xi^{-1}$, we have by $\psi_{2,6}$ that $\sigma'(\xi') \neq c$.

Moreover, since $\mathfrak{A} \models r_{\sigma', c', \xi'}[\tilde{a}a]$, by $\psi_{2,5}$ we have $\mathfrak{A} \models (\sigma'(\xi'))[a\bar{a}]$, contradicting $\mathfrak{A} \models c[a\bar{a}]$. Thus $\xi' = \xi^{-1}$ as required.

This completes the process of finding witnesses for all ℓ -tuples from A , in the course of which we have partially defined a layered structure \mathfrak{B} , such that, however it is completed, $\mathfrak{B} \models \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_i$ for every $i \in I$. Moreover, each of the newly-fixed primitive $(\ell+1)$ -tuples (and their reversals) satisfies β .

All that remains is to complete the construction of \mathfrak{B} by assigning adjacent types to all primitive $(\ell+1)$ -tuples whose adjacent types have not yet been fixed, without violating the condition $\forall \mathbf{x}_{\ell+1} \beta$. Suppose, then $\bar{c} = a\bar{a}b$ is such a primitive $(\ell+1)$ -tuple. By primitiveness, any defect $\langle i, j \rangle$ of \bar{c} is such that $2 < i < j \leq \ell$ and $j-i+1$ is odd. Thus, writing D for the set of defects of \bar{a} , we have that $D \subseteq \mathbf{D}_{\ell-1}^\circ$, and D^+ is exactly the defect set of \bar{c} . By ψ_1 we have $\mathfrak{A} \models \delta_D[\bar{a}]$ and, writing $\zeta = \text{atp}^{\mathfrak{A}^-}[a\bar{a}]$, we have that $\mathfrak{A} \models p_\zeta[\bar{a}]$ by $\psi_{2,0}$. And writing $\eta = \text{atp}^{\mathfrak{A}^-}[\bar{a}b]$, by ψ_4 there exists an adjacent $(\ell+1)$ -type ξ such that $\xi \models \zeta \wedge \hat{\beta} \wedge \eta^+$ and $\partial\xi$ is D^+ -compatible. By the second statement of Lemma 8, then, we can consistently assign $\text{itp}^{\mathfrak{B}}[\bar{c}] = \partial\xi$. Since $\zeta = \text{atp}^{\mathfrak{A}^-}[a\bar{a}]$, $\eta = \text{atp}^{\mathfrak{A}^-}[\bar{a}b]$, and $\xi = \zeta \cup \eta^+ \cup \partial\xi$, we have $\text{atp}^{\mathfrak{B}}[\bar{c}] = \xi$. Additionally, since $\xi \models \hat{\beta}$, both \bar{c} and \bar{c} satisfy the universal requirements β in \mathfrak{B} . Repeated applications of this procedure result in all primitive $(\ell+1)$ -tuples having their adjacent types defined in such a way that $\mathfrak{B} \models \forall \mathbf{x}_{\ell+1} \beta$. Hence $\mathfrak{B} \models \varphi$. \dashv

Taken together, Lemmas 15 and 16 reduce the satisfiability problem for $\mathcal{AF}^{\ell+1}$ to that for \mathcal{AF}^ℓ ($\ell \geq 2$), though with exponential blow-up. We thus obtain the decidability of satisfiability for the whole of \mathcal{AF} . More precisely:

THEOREM 1. *If φ is a satisfiable $\mathcal{AF}^{\ell+1}$ -formula, with $\ell \geq 1$, then φ is satisfied in a structure of size at most $\mathfrak{t}(\ell, \|\varphi\|^{O(1)})$. Hence the satisfiability problem for \mathcal{AF}^k is in $(k-1)$ -NEXPTIME for all $k \geq 2$, and the adjacent fragment is TOWER-complete.*

PROOF. Fix $\ell \geq 2$ and suppose φ is a satisfiable $\mathcal{AF}^{\ell+1}$ -formula over a signature τ . By Lemma 1, we may assume that φ is in normal form. Writing $\varphi_{\ell+1}$ for φ , let φ_ℓ now denote the formula $\psi := \text{acl}(\varphi) \wedge \psi_0 \wedge \dots \wedge \psi_5$ as defined above. Repeating this process, we obtain a sequence of formulas $\varphi_{\ell+1}, \dots, \varphi_2$. By Lemma 15, φ_2 is then also satisfiable. For all k , ($2 \leq k \leq \ell+1$), let φ_k have signature τ_k , and for $k \leq \ell$, consider the construction of φ_k from φ_{k+1} . Since $\sum_{k'=1}^{k+1} |\mathbf{A}_{k'}^{k+1}|$ is bounded by a constant, we see that $\|\text{acl}(\varphi)_{k+1}\|$ is $O(\|\varphi_{k+1}\|)$. In regard to ψ_0, \dots, ψ_5 we may, when considering the adjacent $(k+1)$ -types over τ_{k+1} , disregard all adjacent atoms whose argument sequence is not a substitution instance of some argument sequence \mathbf{x}_{k+1}^g occurring in an atom of φ_{k+1} , as these cannot affect the evaluation of φ_{k+1} . And since $k \leq \ell$, the number of functions from \mathbf{x}_{k+1} to itself is again bounded by a constant, and the number of adjacent $(k+1)$ -atoms over τ_{k+1} that impact the satisfaction of φ_{k+1} is $O(\|\varphi_{k+1}\|)$. Thus, the number of adjacent $(k+1)$ -types over τ_{k+1} that we need to consider is $2^{O(\|\varphi_{k+1}\|)}$. (This cardinality bounds the number of predicates p_ζ as featured in $\psi_{2,0}$ and ψ_4). Keeping the number of $(k+1)$ -types we need to consider in mind recall that a star is a partial function mapping $(k+1)$ -types to a set of colours C . Pick any such star σ and let m be the cardinality of $\text{dom}(\sigma)$. Clearly, m is

at most $\|\varphi_{k+1}\|$ as no tuple needs more witnesses than there are existential requirements. On the other hand, $|C|$ is also bounded by the number of existential requirements in φ_{k+1} and is at most $2(\|\varphi_{k+1}\|^2 + \|\varphi_{k+1}\|) + 1$. Noting that there are at most $|C|^m$ functions over a domain of size m and image C , we have that the total number of stars does not exceed $\sum_{m=0}^{\|\varphi_{k+1}\|} (|C|^m \cdot (2^{O(\|\varphi_{k+1}\|)})^m)$ which is clearly bounded by $2^{\|\varphi_{k+1}\|^{O(1)}}$. Taking these considerations together and noting that the set of defects \mathbf{D}_{k-1}° is bounded by the constant $(k-1)^2$, we have that each of the sentences ψ_0, \dots, ψ_5 contain at most $2^{\|\varphi_{k+1}\|^{O(1)}}$ conjuncts. Some care is needed when calculating the sizes of these conjuncts themselves, as they feature disjunctions over stars as in $\psi_{2,6}$ or that over tailing ends of D^+ -compatible $(k+1)$ -types obtained by processing $\zeta \wedge \hat{\beta}$ for some set of defects $D \subseteq \mathbf{D}_{k-1}^\circ$ as in ψ_4 . However, in the first case, the disjunction ranges over a maximum of $2^{\|\varphi_{k+1}\|^{O(1)}}$ predicates $s_{\sigma'}$, and in the second case the formulas are, in effect, in disjunctive normal form over atoms contained in φ_{k+1} , and hence have cardinality $2^{O(\|\varphi_{k+1}\|)}$. Hence, $\|\varphi_k\|$ is $2^{\|\varphi_{k+1}\|^{O(1)}}$. By an easy induction, $\|\varphi_2\|$ is then $\mathfrak{t}(\ell-1, \|\varphi_{\ell+1}\|^{O(1)})$.

Since φ_2 is a sentence in the two-variable fragment of first-order logic, we have, by [16], that φ_2 has a model \mathfrak{A}_2 of cardinality $2^{O(\|\varphi_2\|)}$. Moreover, by Lemma 16, each of the formulas φ_{k+1} ($2 \leq k \leq \ell$) have a model over the same domain as \mathfrak{A}_2 . Since $|\tau_k| \geq |\tau_{k+1}|$, we conclude that the expansion of \mathfrak{A}_2 (and subsequent models) into models $\mathfrak{A}_3, \dots, \mathfrak{A}_{\ell+1}$ of $\varphi_3, \dots, \varphi_{\ell+1}$ remains bounded by $2^{O(\|\varphi_2\|)}$ and thus is of size $\mathfrak{t}(\ell, \|\varphi_{\ell+1}\|^{O(1)})$ as required. \dashv

We note that one can do better if the equality predicate is disallowed. As shown in [8, Lemma 4.6], \mathcal{AF}^3 -sentences have a model exponential size in regard to the input. Thus, by reducing variables of $\varphi_{\ell+1}$ as in Theorem 1 until an equisatisfiable \mathcal{AF}^3 -sentence φ_3 is reached, we have that φ_3 (and thus also $\varphi_{\ell+1}$) has a model of size $\mathfrak{t}(\ell-1, \|\varphi_{\ell+1}\|^{O(1)})$ thus establishing the following:

THEOREM 2. *The satisfiability problem for the sub-fragment of \mathcal{AF}^k without equality is in $(k-2)$ -NEXPTIME for all $k \geq 3$.*

It may not have escaped the reader's attention that the equality predicate barely features in the proof of Theorem 1. This is because the reduction from the satisfiability problem for $\mathcal{AF}^{\ell+1}$ to that for \mathcal{AF}^ℓ effected in Lemmas 15 and 16 does not interfere with predicates of arity less than $(\ell+1) \geq 3$, and does not interfere with the domain of quantification. The special status of the equality predicate is thus, in a sense, pushed back into the two-variable fragment.

This observation suggests a generalization. Recent decades have witnessed concerted attempts to investigate the decidability of the satisfiability problem for FO^2 over various classes of structures, where certain distinguished predicates are required to be interpreted as relations subject to various restrictions, e.g. as linear orders [25, 37], trees [13, 6], equivalences [25, 23] and more (see [24] for a survey). If such an extension of FO^2 is decidable for (finite) satisfiability, then the reduction outlined in this section applies to the corresponding extension of \mathcal{AF} —it being understood that the distinguished predicates in question have arity at most 2. On the other hand, if the extension of FO^2 is undecidable for (finite)

satisfiability, then—anticipating Theorem 4, which states that \mathcal{AF} subsumes FO^2 —this undecidability result likewise immediately transfers to \mathcal{AF} .

§6. The Guarded Subfragment. We next shift our attention to the *guarded adjacent fragment*, denoted \mathcal{GA} , defined as the intersection of \mathcal{AF} with the guarded fragment, \mathcal{GF} [1, Sec. 4.1]. In \mathcal{GF} , quantification is relativized by atoms, i.e. all quantification takes the form $\forall \bar{x}(\alpha \rightarrow \psi)$ and $\exists \bar{x}(\alpha \wedge \psi)$, where α (a *guard*) is an atom featuring all the variables in \bar{x} and all the free variables of ψ . The satisfiability problem for \mathcal{GF} is 2EXPTIME -complete [15, Thm. 4.4]. We show that the satisfiability problem for \mathcal{GA} remains 2EXPTIME -complete, and thus is as hard as for full \mathcal{GF} . This contrasts with the EXPTIME -completeness of the k -variable guarded fragment \mathcal{GF}^k [15, Cor. 4.6] and of the guarded forward fragment [5, Thm. 4]. Our proof follows the same reduction strategy as the 2EXPTIME -hardness proof for \mathcal{GF} by E. Grädel [15, Thm. 4.4]. Because we are working in the guarded *adjacent* fragment, however, Grädel’s reduction is not directly available, and some new techniques are required. To aid readability, we employ the variable names w, x, y, z, \dots in adjacent formulas, rather than the official $x_1, x_2, x_3, x_4, \dots$. The reader may easily check that all formulas in question are indeed in \mathcal{AF} modulo variable renaming.

6.1. Generating Words. We start with a combinatorial exercise concerning the generation of words by the recurrent application of certain adjacent functions. Let $m \in \mathbb{N}$ and consider the adjacent functions $\lambda_1, \lambda_2, \lambda_3: [1, m+2] \rightarrow [1, m+1]$ defined by the following courses of values $(\lambda_i(1) \ \dots \ \lambda_i(m+2))$:

$$\begin{aligned} \lambda_1: & \quad (1 \ 2 \ 2 \ 3 \ 4 \ \dots \ m+1) \\ \lambda_2: & \quad (1 \ 2 \ 1 \ 2 \ 3 \ \dots \ m) \\ \lambda_3: & \quad (1 \ 2 \ 3 \ 3 \ 4 \ \dots \ m+1). \end{aligned}$$

Using the imagery of Sec. 3, we can think of these functions as describing strolls on a word of length $(m+1)$, in each case starting at the left-most position, and proceeding generally rightwards: λ_1 pauses at the second time step before continuing; λ_2 returns to the beginning after the second time step, but then resumes its rightward journey (without quite reaching the end); and λ_3 pauses at the third time step before continuing.

We show that repeated application of these functions to the bit-string 011^m yields the whole of the (exponentially large) language $01\{0, 1\}^m$.

LEMMA 17. *Let $W_0 \subseteq \{0, 1\}^*$ contain 011^m and $W_i := W_{i-1} \cup \{\bar{w}^{\lambda_1}, \bar{w}^{\lambda_2}, \bar{w}^{\lambda_3} \mid \bar{w} \in W_{i-1}\}$. Setting $W := \bigcup_{i \geq 0} W_i$, we have $01\{0, 1\}^m \subseteq W$.*

PROOF. We establish by induction on $i \in [0, m]$ that, for any word $\bar{c} \in \{0, 1\}^i$, the word $01\bar{c}1^{m-i}$ is in W ; the case $i = m$ then yields the statement of the lemma. The base case, $i = 0$ follows from the assumption that W_0 contains 011^m . Suppose now $i > 0$. We show that, for any $x \in \{0, 1\}$, the word $\bar{u} = 01x\bar{c}1^{m-i-1}$ is in W . Letting c_1 be the first character of \bar{c} , and writing $\bar{c} = c_1\bar{d}$, it follows by inductive hypothesis that the words $\bar{v} = 01\bar{c}1^{m-i-1}1$ and $\bar{w} = 01\bar{d}1^{m-i-1}11$ are in W . Consider cases: (i) if $x = 1$ then $\bar{u} = \bar{v}^{\lambda_1}$, (ii) if both $x = 0$ and $c_1 = 0$ then $\bar{u} = \bar{v}^{\lambda_3}$, and otherwise, (iii) $x = 0$, $c_1 = 1$ and $\bar{u} = \bar{w}^{\lambda_2}$. Thus $\bar{u} \in W$. \dashv

We now apply Lemma 17 to \mathcal{GA} . For $m \geq 0$, let G_m be an $(m+2)$ -ary predicate. We proceed to write, for any binary predicate p , a \mathcal{GA} -sentence $\overrightarrow{\text{allSeq}}_m(p)$ ensuring that, if p is satisfied by a pair of objects, say ab , then G_m is satisfied by any $(m+2)$ -tuple $ab\bar{w}$ for $\bar{w} \in \{a, b\}^m$. By Lemma 17, it suffices so take $\overrightarrow{\text{allSeq}}_m(p)$ to be

$$\forall xy \left(p(xy) \rightarrow G_m(xy \underbrace{y \cdots y}_m) \right) \wedge \bigwedge_{i \in [1,3]} \forall \mathbf{u}_{m+2} \left(G_m(\mathbf{u}_{m+2}) \rightarrow G_m(\mathbf{u}_{m+2}^{\lambda_i}) \right).$$

Thus, if $\mathfrak{A} \models \overrightarrow{\text{allSeq}}_m(p)$, and $\mathfrak{A} \models p[ab]$, we may freely quantify over words of length $(m+2)$ over the alphabet $\{a, b\}$ as long as they have the prefix ab , since G_m can always be used as a guard.

We shall require a ‘mirrored’ version of the above device, this time involving a pair of 2-element alphabets and a pair of words of length m over these alphabets. For $m \geq 0$, let F_m be a $(2m+4)$ -ary predicate. We proceed to write, for any quaternary predicate r , a \mathcal{GA} -sentence $\overleftarrow{\text{allSeq}}_m(r)$ ensuring that, if r is satisfied by a quartet of objects $abcd$, then, for any m -tuple \bar{u} over the alphabet $\{a, b\}$, and any m -tuple \bar{v} over the alphabet $\{c, d\}$, the predicate F_m is satisfied by $\bar{u}abcd\bar{v}$. Let λ_0 denote the identity function on $[1, m+2]$, and take $\lambda_1, \lambda_2, \lambda_3$ as defined above. In addition, for any word \bar{w} , we write \bar{w}^R for the reversal of \bar{w} . (Thus, $\bar{w}^R = \bar{w}$; but the new notation is more readable in what follows.) By two applications of Lemma 17, it suffices so take $\overleftarrow{\text{allSeq}}_m(r)$ to be:

$$\forall yxzt \left(r(yxzt) \rightarrow F_m(\underbrace{y \cdots y}_m yxzt \underbrace{t \cdots t}_m) \right) \wedge \bigwedge_{i,j \in [0,3]} \forall \mathbf{u}_{m+2}^R \mathbf{v}_{m+2} \left(F_m(\mathbf{u}_{m+2}^R \mathbf{v}_{m+2}) \rightarrow F_m((\mathbf{u}_{m+2}^{\lambda_i})^R \mathbf{v}_{m+2}^{\lambda_j}) \right).$$

Again, if $\mathfrak{A} \models \overleftarrow{\text{allSeq}}_m(r)$, and $\mathfrak{A} \models r[abcd]$, we may freely quantify over words of the language $\{a, b\}^m b a c d \{c, d\}^m$ in \mathcal{GA} , since F_m can always be used as a guard. Note that the variables of $\mathbf{u}_{m+2} = u_1 \cdots u_{m+2}$ are quantified above in ‘reverse order’. This ensures, after renaming, the adjacency of the formula $\overleftarrow{\text{allSeq}}_m(r)$.

6.2. ATMs. An *alternating Turing machine* (ATM) [12] is a tuple $\mathcal{M} := \langle Q, \Sigma, q_0, \Delta \rangle$, where Q is a non-empty finite set (the *states* of \mathcal{M}), Σ a non-empty finite alphabet, q_0 an element of Q (the *initial state*), and Δ a set *transitions* δ , defined presently. We imagine \mathcal{M} to operate on an 1-way infinite tape by means of a read-write head as usual, but we take Q to be partitioned into the sets Q_{\exists} (*existential states*), Q_{\forall} (*universal states*) and $\{q_a, q_r\}$ (the *accepting* and *rejecting state*, respectively). Writing Σ' for the alphabet Σ augmented with the blank cell symbol ‘ \sqcup ’, we define a *transition* $\delta \in \Delta$ to be a relation

$$\delta \in Q \times \Sigma' \times Q \times \Sigma' \times \{-1, 0, 1\}.$$

The transition $\delta = \langle q, s, q', s', k' \rangle$ is enabled when the machine is in state q and the read-write head is positioned over a tape square containing the letter s . On execution of δ , the current tape square is overwritten with s' , the head is moved by k' squares, and the current state is updated to q' . We denote the set of $\delta \in \Delta$ enabled by state q and letter s by $\Delta(q, s)$.

A *configuration* \mathcal{C} of \mathcal{M} is a triple $\langle q, \omega, h \rangle$, where $q \in Q$, ω is an infinite word over Σ' and h a non-negative integer. We read the triple \mathcal{C} as stating that the machine is in state q , the tape contents are given by ω , and the head is situated over the h th tape square (counting from 0). If \bar{w}_0 is a finite word over Σ representing the input to the machine, the *initial configuration* is $\langle q_0, \bar{w}_0 \sqcup^*, 0 \rangle$, where \sqcup^* represents an infinite series of blanks. The *successors* of \mathcal{C} are defined in the usual way via transitions in Δ . A *halting* configuration is one which is in state q_a or q_r . We assume that halting configurations have no enabled transitions, and non-halting configurations always have at least one enabled transition.

We shall be interested in the case where every computation of \mathcal{M} (understood as a sequence of enabled transitions starting in the initial configuration) is of finite length, so that there is a function f such that, when \mathcal{M} runs on input w_0 , the read-write head never reaches positions beyond $f(|w_0|)-1$. In that case, we may as well take a configuration to have the form $\langle q_0, \bar{w}, h \rangle$ where \bar{w} is a (finite) word over Σ' of length $f(|w_0|)-1$ and h is an integer in the range $[0, f(|w_0|)-1]$. The notions of acceptance and rejection may then be defined as follows: a halting configuration is *accepting* if it is in state q_a ; an existential configuration is *accepting* if it has an accepting successor; a universal configuration is *accepting* if all its successors are accepting; a configuration which is not accepting is *rejecting*. We take \mathcal{M} to accept \bar{w}_0 if the initial configuration is accepting.

We witness acceptance of an input \bar{w}_0 by \mathcal{M} using an *acceptance tree* \mathcal{T} . This is a finite tree with vertices labelled by (accepting) configurations. The root is labelled by an *initial configuration*; and for any vertex labelled with a particular configuration, its children are labelled with the results of executing enabled transitions in that configuration. Vertices labelled with existential configurations have *at least one* child corresponding to an enabled transition; those labelled with universal configurations have a child corresponding to *every* enabled transition; the leaves of the tree are labelled with accepting configurations.

We are interested in the case where the function f bounding the space required by \mathcal{M} is of the form $f(n) = 2^n$. Thus, \mathcal{M} accesses at most $2^{|\bar{w}_0|}$ tape squares in the course of any computation on input \bar{w}_0 . We now fix such a machine \mathcal{M} , and show how, for a given input \bar{w}_0 , we can manufacture a \mathcal{GA} -formula $\varphi_{\mathcal{M}, \bar{w}_0}$ satisfiable if and only if \mathcal{M} accepts \bar{w}_0 . The computation of $\varphi_{\mathcal{M}, \bar{w}_0}$ runs in polynomial time (in fact, in logarithmic space) as a function of $|\bar{w}_0|$. Since there are problems in $\text{ASPACE}(2^n)$ that are complete for AEXPSPACE , we can thereby reduce any problem in AEXPSPACE to the satisfiability problem for \mathcal{GA} . Hence the latter problem is AEXPSPACE -hard. Using the well-known equation $\text{AEXPSPACE} = 2\text{EXPTIME}$, this achieves our goal.

6.3. Encoding numbers. In the sequel, we will consider structures \mathfrak{A} interpreting a unary predicate O . Whenever $\mathfrak{A} \models \alpha[ab]$, where $\alpha(xy)$ is the formula $\neg O(x) \wedge O(y)$, we say that a and b *act as zero and unit bits* (O for “One”), and for any word \bar{u} over $\{a, b\}$, we write $\text{val}^{\mathfrak{A}}(\bar{u})$ to denote the integer represented by \bar{u} , considered as a bit-string, with a standing for 0 and b for 1 (most significant bit first). Notice that there may be other elements, say c and d , such that $\mathfrak{A} \models \alpha[cd]$, in which case may write $\text{val}^{\mathfrak{A}}(\bar{w})$ for the integer represented by any word \bar{w} over $\{c, d\}$. Clearly, the \mathcal{GA} -formula $\text{EQ}(\mathbf{u}_n, \mathbf{v}_n) := \bigwedge_{i=1}^n O(u_i) \leftrightarrow O(v_i)$

satisfies $\mathfrak{A} \models \text{EQ}[\bar{c}, \bar{d}]$ if and only if $\text{VAL}^{\mathfrak{A}}(\bar{c}) = \text{VAL}^{\mathfrak{A}}(\bar{d})$. Other arithmetical properties can be expressed similarly. In particular, for each $k \in \{-1, +1\}$, we may easily write a \mathcal{GA} -formula $\text{EQ}(\mathbf{u}_n, \mathbf{v}_n + k)$ satisfying the following property:

$$\mathfrak{A} \models \text{EQ}[\bar{c}, \bar{d} + k] \text{ iff } \text{VAL}^{\mathfrak{A}}(\bar{c}) = \text{VAL}^{\mathfrak{A}}(\bar{d}) + k.$$

The details are routine (see [8, p. 15]). Observe that the formula $\text{EQ}(\bar{x}, \bar{y})$ is not atomic! Indeed, all the predicates appearing in this formula are unary; thus, this formula may appear in adjacent formulas whatever the order of quantification among the variables \bar{x} and \bar{y} . Similar remarks apply to $\text{EQ}(\bar{x}, \bar{y} + k)$.

6.4. Encoding configurations. We proceed to describe a method of encoding, within a certain class of structures, configurations of an ATM $\mathcal{M} = \langle Q, \Sigma, q_0, \Delta \rangle$ that never accesses more than $f(n) = 2^n$ tape squares on an input of size n . Recall that we agreed to regard configurations in this case as triples $\langle q, \bar{w}, h \rangle$, where $q \in Q$, \bar{w} is a word over Σ' of length $N = 2^n$, and $0 \leq h < N$. In doing so, we take the various states $q \in Q$ to be *binary* predicates: a configuration is then represented by an ordered pair of (distinct) elements ab satisfying any of these predicates. If $\mathfrak{A} \models q[ab]$, we take it that the represented configuration has state q . Since we shall want words over the alphabet $\{a, b\}$ to represent integers, we shall insist that, in this case, ab satisfies the formula α defined in Sec. 6.3. This we ensure by writing the \mathcal{GA} -sentence $\bigwedge_{q \in Q} \forall xy (q(xy) \rightarrow \neg O(x) \wedge O(y))$. (Thus, we think of a as a 0 and b as a 1.) To represent further aspects of the configuration ab , we treat each symbol $s \in \Sigma'$ as an n -ary predicate, and additionally employ an n -ary predicate H . Specifically, for any $\bar{w} \in \{a, b\}^n$, we read $\mathfrak{A} \models H[\bar{w}]$ as “the head of the configuration represented by ab is at position $\text{VAL}^{\mathfrak{A}}(\bar{w})$ ”, and we read $\mathfrak{A} \models s[\bar{w}]$ as “the tape square $\text{VAL}^{\mathfrak{A}}(\bar{w})$ of the configuration represented by ab contains the symbol s ”. Of course, these interpretations are only meaningful if:

1. there is at most one string in $\{a, b\}^n$ satisfying H and thus encoding the head position;
2. each bit-string over $\{a, b\}^n$ satisfies at most one predicate $s \in \Sigma'$, thus ensuring that a tape cell contains at most one symbol; and
3. ab satisfies at most one $q \in Q$, thus ensuring that the configuration is in at most one state.

For any state $q \in Q$ we construct a \mathcal{GA} -formula $\psi_q := \psi_{q,1} \wedge \psi_{q,2} \wedge \psi_{q,3}$ enforcing these conditions for any configuration whose state is q . The first conjunct, $\psi_{q,1}$, may be given as follows:

$$\psi_{q,1} := \overrightarrow{\text{allSeq}}_{2^n}(q) \wedge \forall xy \mathbf{u}_n \mathbf{v}_n \left(G_{2^n}(xy \mathbf{u}_n \mathbf{v}_n) \rightarrow \left((H(\mathbf{u}_n) \wedge H(\mathbf{v}_n)) \rightarrow \text{EQ}(\mathbf{u}_n, \mathbf{v}_n) \right) \right).$$

Note that this sentence is (adjacent and) guarded, with the atom $G_{2^n}(xy \mathbf{u}_n \mathbf{v}_n)$ acting as a guard. Of course, thanks to the conjunct $\overrightarrow{\text{allSeq}}_{2^n}(q)$, this guard is, as it were, *semantically inert*, since, assuming that q is satisfied by ab , G_{2^n} is satisfied by *every* word of $ab\{a, b\}^{2^n}$. The remaining conjuncts $\psi_{q,2}$ and $\psi_{q,3}$ can be easily formulated.

6.5. Encoding instances of transitions. For each transition $\delta \in \Delta$ we employ a quaternary predicate E_δ and read $\mathfrak{A} \models E_\delta[abcd]$ as “the configuration encoded by ab enables a transition δ thus producing the configuration encoded by cd ”. (Do note the reversal of ab in $E_\delta[abcd]$!) We call $abcd$ a δ -transition instance in \mathfrak{A} with ab corresponding to the predecessor configuration and cd to the successor configuration. To make sure that each δ -transition instance is indeed a result of transitioning via $\delta = \langle q, s, q', s', k' \rangle$ we write the following in \mathcal{GA} :

1. the successor configuration is in state q' ;
2. the head of the successor configuration is moved by k' relative to the predecessor configurations head;
3. the h -th tape cell on the successor configuration is occupied by s' , where h is the position of the predecessor's head;
4. all tape cells that the predecessor's head does not point to are inherited by the successor.

For any transition $\delta \in \Delta$, we construct a \mathcal{GA} -formula $\psi_\delta := \psi_{\delta,1} \wedge \cdots \wedge \psi_{\delta,4}$ enforcing these conditions. We write $\psi_{\delta,4}$ in detail as an example; the preceding three conjuncts are handled similarly or more easily:

$$\psi_{\delta,4} := \overleftarrow{\text{allSeq}}_n(E_\delta) \wedge \forall \mathbf{u}_n yxzt \mathbf{v}_n \left(F_n(\mathbf{u}_n yxzt \mathbf{v}_n) \rightarrow \left((E_\delta(yxzt) \wedge \neg H(\mathbf{u}_n) \wedge \text{EQ}(\mathbf{u}_n, \mathbf{v}_n)) \rightarrow \bigwedge_{s \in \Sigma} (s_{\mathbf{a}}(\mathbf{u}_n) \leftrightarrow s_{\mathbf{a}}(\mathbf{v}_n)) \right) \right).$$

Suppose now that $\mathfrak{A} \models \psi_\delta$, and moreover, that $\mathfrak{A} \models E_\delta[abcd]$ for some $abcd \in A^4$, where ab and cd both encode configurations. By $\overleftarrow{\text{allSeq}}_n(E_\delta)$, we have that $\mathfrak{A} \models F_n[\bar{u}abcd\bar{v}]$ for all $\bar{u} \in \{a, b\}^n$ and $\bar{v} \in \{c, d\}^n$. By picking any $0 \leq i < 2^n$ that is not the head position of the configuration encoded by ab , and $\bar{u} \in \{a, b\}^n$, $\bar{v} \in \{c, d\}^n$ with $\text{VAL}^{\mathfrak{A}}(\bar{u}) = \text{VAL}^{\mathfrak{A}}(\bar{v}) = i$, we are guaranteed that, for each symbol $s \in \Sigma$, $\mathfrak{A} \models s[\bar{u}]$ if and only if $\mathfrak{A} \models s[\bar{v}]$.

6.6. Encoding acceptance trees. The last step in our reduction is to write a formula whose models contain configurations arranged as an acceptance tree witnessing the fact that \mathcal{M} accepts some input $w_0 = s_1 \cdots s_n$. Recall that, the root of this tree is labelled with the initial configuration, namely $\mathcal{C} = \langle q_0, \bar{w}_0 \sqcup^\ell, 0 \rangle$, where $\ell = 2^n - |\bar{w}_0|$. We ensure the existence of such a root configuration in a structure with the following \mathcal{GA} -sentence:

$$\overrightarrow{\text{allSeq}}_n(q_0) \wedge \exists xy \left(q_0(xy) \wedge H(\text{“0”}) \wedge \bigwedge_{i=0}^{i < |\bar{w}_0|} s_{i+1}(\text{“}i\text{”}) \wedge \forall \bar{u} \left(G(xy\bar{u}) \rightarrow \left(\left(\bigwedge_{i=0}^{i < |\bar{w}_0|} \neg \text{EQ}(\text{“}i\text{”, } \bar{u}) \right) \rightarrow \sqcup(\bar{u}) \right) \right) \right),$$

where “ i ” is the binary encoding of i using x as a zero bit and y as a unit bit. We ensure the existence of successor configurations required by existential configurations as follows. Suppose that q is an existential state, and s a symbol. The following sentence ensures that any configuration in state $q \in Q_\exists$ with the head

reading symbol $s \in \Sigma'$ has a child in the acceptance tree:

$$\overrightarrow{\text{allSeq}}_n(q) \wedge \forall \tilde{\mathbf{u}}_n y x \left(G_n(xy \mathbf{u}_n) \rightarrow \right. \\ \left. ((q(xy) \wedge H(\mathbf{u}_n) \wedge s(\mathbf{u}_n)) \rightarrow \bigvee_{\delta \in \Delta(q,s)} \exists zt E_\delta(yxzt)) \right).$$

In case q is universal, the disjunction over $\Delta(q, s)$ is replaced by a conjunction. Lastly, we write the sentence $\neg \exists xy q_r(xy)$ to ensure that a rejecting configuration is never encoded by any pair of elements in any structure.

For any input string w_0 , let $\varphi_{\mathcal{M}, \bar{w}_0}$ be the conjunction of all the \mathcal{GA} -sentences given above. We remark that $\varphi_{\mathcal{M}, \bar{w}_0}$ does not feature the equality predicate. Bearing in mind that \mathcal{M} is guaranteed to terminate on \bar{w}_0 in a finite number of steps, accessing no more than $2^{|\bar{w}_0|}$ tape squares, we see that any model of $\varphi_{\mathcal{M}, \bar{w}_0}$ embeds an acceptance tree for \mathcal{M} on input \bar{w}_0 . Conversely, any acceptance trees for \mathcal{M} on input \bar{w}_0 can be expanded to a model of $\varphi_{\mathcal{M}, \bar{w}_0}$ by interpreting the relevant predicates as suggested above. We conclude:

THEOREM 3. *The finite and general satisfiability problem for the (equality-free) guarded adjacent fragment of first-order logic is 2EXPTIME-hard.*

§7. Extending the adjacent fragment. In this final section, we consider the prospects for extending the logic \mathcal{AF} within the family of argument-sequence fragments. As a preliminary, we compare the expressive power of \mathcal{AF} to that of the two-variable fragment of first-order logic.

Denote by FO^2 the two-variable fragment of first-order logic, namely, the set of first-order formulas over a purely relational signature (i.e. no individual constants or function symbols) in which the only logical variables occurring are x_1 and x_2 . We allow the equality predicate in FO^2 .

THEOREM 4. *Every FO^2 -formula is logically equivalent to an \mathcal{AF} -formula. Conversely, every \mathcal{AF} -sentence featuring predicates of arity at most two is logically equivalent to an FO^2 -sentence.*

PROOF. Let Q stand for either of the quantifiers \forall or \exists . For the first statement of the theorem, suppose that φ is an FO^2 -formula. Let k be the highest index of any free variable of φ , and $k = 0$ if φ is a sentence. (Thus, $0 \leq k \leq 2$.) We claim that φ is logically equivalent to an $\mathcal{AF}^{[k]}$ -formula, proceeding by structural induction on φ . For the base case, φ is an atom of FO^2 , and therefore certainly in $\mathcal{AF}^{[k]}$. The case of Boolean operators is routine (bearing in mind that $\mathcal{AF}^{[0]} \subseteq \mathcal{AF}^{[1]} \subseteq \mathcal{AF}^{[2]}$). Now consider the case where φ has the form $Qx_1 \psi$. If x_1 does not occur free in ψ , then φ and ψ are logically equivalent, and the result follows immediately by inductive hypothesis. If x_1 does occur free in ψ , but x_2 does not, then, by inductive hypothesis, let $\psi' \in \mathcal{AF}^{[1]}$ be logically equivalent to ψ . Then $\varphi' = Qx_1 \psi' \in \mathcal{AF}^{[0]}$ is logically equivalent to φ as required. Otherwise, both x_1 and x_2 occur free in ψ : let φ^* be the result of transposing the variables x_1 and x_2 everywhere in φ , and similarly for ψ . Thus φ is satisfied in a structure \mathfrak{A} under an assignment in which $x_2 \leftarrow a$ if and only if φ^* is satisfied in \mathfrak{A} under an assignment in which $x_1 \leftarrow a$. Moreover, $\varphi^* = Qx_2 \psi^*$, with x_2 free in ψ^* . By

inductive hypothesis, ψ^* is equivalent to a formula $\psi' \in \mathcal{AF}^{[2]}$, and hence φ^* is equivalent to $Qx_2\psi' \in \mathcal{AF}^{[1]}$. Now let $\varphi' \in \mathcal{AF}^{[2]}$ be the result of incrementing the indices of all variables (free or bound) in $Qx_2\psi'$. Thus φ' is satisfied in \mathfrak{A} under an assignment in which $x_2 \leftarrow a$ if and only if φ^* is satisfied in \mathfrak{A} under any assignment in which $x_1 \leftarrow a$. That is to say, φ' is logically equivalent to φ as required. The remaining case is where φ is of the form $Qx_2\psi$. If x_2 does not occur free in ψ , then φ and ψ are logically equivalent, and the result follows immediately by inductive hypothesis. If x_1 does not occur free in φ (i.e. φ is a sentence), then φ is logically equivalent to $\varphi^* = Qx_1\psi^*$, and by the previous case, φ^* is equivalent to a formula of $\mathcal{AF}^{[0]}$. Otherwise, if both x_1 and x_2 occur free in ψ , by inductive hypothesis, ψ is logically equivalent to a formula $\psi' \in \mathcal{AF}^{[2]}$, whence φ is logically equivalent to $Qx_2\psi' \in \mathcal{AF}^{[1]}$. Since x_1 is free in φ , this is what is required, completing the induction.

Turning to the second statement of the theorem, denote by FO^{2*} the fragment of first-order logic consisting of those formulas with the property that no sub-formula contains more than two free variables. A simple variable re-naming procedure shows that any sentence of FO^{2*} is logically equivalent to a sentence of FO^2 . We prove the following claim: if φ is a formula of \mathcal{AF}^k featuring predicates of arity at most 2, then there exists a formula φ' such that: (i) φ and φ' have the same free variables and are logically equivalent; and (ii) φ' is a Boolean combination of atomic $\mathcal{AF}^{[k]}$ -formulas and FO^{2*} -formulas having at most one free variable. Since every sentence of \mathcal{AF} is (by adding vacuous quantification if necessary) an $\mathcal{AF}^{[0]}$ -formula, and since any FO^{2*} -sentence is logically equivalent to an FO^2 -sentence, putting $k = 0$ yields the desired result. Observe that, since all predicates have arity at most 2, atomic formulas can contain at most two variables; and since φ is in \mathcal{AF} , any two distinct free variables occurring in an atomic formula must have indices differing by exactly one. To avoid special cases in the following proof, we assume that x_0 is a variable, though this variable will never actually appear in any formulas. We prove the claim by structural induction on φ . If φ is atomic, the claim is immediate. Likewise, if φ is $\neg\psi$ or $\psi_1 \circ \psi_2$, where \circ is a Boolean connective, it follows immediately by inductive hypothesis.

Suppose then φ has the form $\forall x_{k'}\psi$. By the formation rules for \mathcal{AF} , we have $\psi \in \mathcal{AF}^{[k']}$. Let ψ' be the formula guaranteed by the inductive hypothesis applied to ψ . Using standard Boolean identities, we may re-write the Boolean combination ψ' in conjunctive normal form, say as $\chi := \chi_1 \wedge \cdots \wedge \chi_\ell$, where each χ_i is a disjunction, say $\theta_{i,1} \vee \cdots \vee \theta_{i,m_i}$, and each $\theta_{i,j}$ is either a literal in $\mathcal{AF}^{[k']}$ or an FO^{2*} -formula having at most one free variable. Thus, φ is logically equivalent to the conjunction $\forall x_{k'}\chi_1 \wedge \cdots \wedge \forall x_{k'}\chi_\ell$. Fixing one of these conjuncts, say $\forall x_{k'}\chi_i$, consider the various disjuncts of χ_i . Let γ_i be the disjunction of those disjuncts $\theta_{i,j}$ in which $x_{k'}$ does not occur free, and let δ_i be the disjunction of the rest. Thus, $\forall x_{k'}\chi_i$ is logically equivalent to $\gamma_i \vee \forall x_{k'}\delta_i$. Clearly, γ_i is a Boolean combination of atomic $\mathcal{AF}^{[k']}$ -formulas and FO^{2*} -formulas with at most one free variable. Moreover, the free variables of γ_i are confined to \mathbf{x}_k , since the free variables of φ are confined to \mathbf{x}_k as well. On the other hand, if $\theta_{i,j}$ occurs in δ_i , then either its free variables are included in $\{x_{k'}\}$, or it is a literal with free variables exactly $\{x_{k'-1}, x_{k'}\}$. Hence, $\forall x_{k'}\delta_i$ is an FO^{2*} -formula with free

variables included in $\{x_{k'-1}\}$. Moreover, if $x_{k'-1}$ actually occurs free in $\theta_{i,j}$, we must have $k'-1 \leq k$, since the free variables of φ are confined to \mathbf{x}_k . Thus, χ and hence φ , is logically equivalent to a Boolean combination of formulas of the required forms.

Suppose finally that φ has the form $\exists x_{k'}\psi$. Then we proceed as in the previous case, but use disjunctive normal form rather than conjunctive normal form.

This completes the induction, and hence the proof. \dashv

Now let us return to the question of extending the adjacent fragment. The adjacent fragment is defined by restricting the permitted argument sequences appearing in atoms. More precisely, within the régime of index-normal formulas, we insist that, in contexts where the variable x_k is available for quantification, all atoms have the form $p(\mathbf{x}_k^f)$, where p is an m -ary predicate, and $f: [1, m] \rightarrow [1, k]$ is an adjacent function. It is natural to ask whether the adjacency restriction might be relaxed without compromising the decidability of satisfiability. As we now show, the answer must be no.

Let $f: [1, m] \rightarrow [1, k]$ be a non-adjacent function, with $m, k \geq 2$: that is, there is some j ($1 \leq j < m$) for which $|f(j+1) - f(j)| \geq 2$. Denote by \mathcal{AF}^f the extension of \mathcal{AF} obtained by allowing atoms of the form $p(\mathbf{x}_k^f)$ in clause (1) of the definition of $\mathcal{AF}^{[k]}$ on page 4. We show that \mathcal{AF}^f is expressive enough to state that a given binary relation is transitive.

LEMMA 18. *Let $f: [1, m] \rightarrow [1, k]$ be a non-adjacent function, T a binary predicate, and Q an m -ary predicate. There exists a formula φ_T of \mathcal{AF}^f , such that: (i) if $\mathfrak{A} \models \varphi_T$ then $T^{\mathfrak{A}}$ is transitive, and (ii) any structure \mathfrak{A} interpreting T as a transitive relation, but not interpreting Q , can be expanded to a model of φ_T .*

PROOF. Since f is non-adjacent, fix an index $j \in [1, m-1]$ for which $|f(j+1) - f(j)| \geq 2$. We may assume without loss of generality that $f(j+1) > f(j)$, as the proof for the other case is obtained by swapping all occurrences of j and $j+1$. Notice that, in this case, we must have $f(j) < f(j+1) - 1$ and hence $f(j) < k - 1$. Define $\varphi_T := \forall \mathbf{x}_m \varphi_T^1 \wedge \forall \mathbf{x}_k \varphi_T^2$, where

$$\begin{aligned} \varphi_T^1 &:= (Q(\mathbf{x}_m) \rightarrow T(x_j, x_{j+1})) \\ \varphi_T^2 &:= \left(T(x_{f(j)}, x_{f(j)+1}) \wedge T(x_{f(j)+1}, x_{f(j)+2}) \wedge \right. \\ &\quad \left. \bigwedge_{i=f(j)+2}^{f(j+1)-1} x_i = x_{i+1} \right) \rightarrow Q(\mathbf{x}_k^f). \end{aligned}$$

To establish (i), we take any $\mathfrak{A} \models \varphi_T$ and any elements a, b, c of \mathfrak{A} with $(a, b) \in T^{\mathfrak{A}}$ and $(b, c) \in T^{\mathfrak{A}}$. We must show that $(a, c) \in T^{\mathfrak{A}}$. Let $\vec{d} = d_1 \cdots d_k$ be a tuple with $d_{f(j)} = a$, $d_{f(j)+1} = b$, and $d_i = c$ for all i except $f(j)$ and $f(j)+1$. Thus, $\mathfrak{A} \models T[d_{f(j)}, d_{f(j)+1}]$, $\mathfrak{A} \models T[d_{f(j)+1}, d_{f(j)+2}]$, and $d_{f(j)+2} = \cdots = d_{f(j+1)}$. Since $\mathfrak{A} \models \varphi_T^2[\vec{d}]$, we have $\mathfrak{A} \models Q[\vec{d}^f]$. But the j th position of \vec{d}^f is occupied by $d_{f(j)} = a$, and the $(j+1)$ th position is occupied by $d_{f(j)+1} = c$; and since $\mathfrak{A} \models \varphi_T^1[\vec{d}^f]$, we have $(a, c) \in T^{\mathfrak{A}}$, as required.

To establish (ii), suppose \mathfrak{A} interprets T as a transitive relation. We expand \mathfrak{A} to \mathfrak{A}^+ by fixing the interpretation of Q to be the set of all m -tuples \vec{a} such that

$(a_j, a_{j+1}) \in T^{\mathfrak{A}}$. It is immediate that $\mathfrak{A}^+ \models \forall \mathbf{x}_m \varphi_T^1$. Now take any k -tuple \bar{c} satisfying the antecedent of φ_T^2 in \mathfrak{A}^+ . Thus, $(c_{f(j)}, c_{f(j)+1})$ and $(c_{f(j)+1}, c_{f(j)+2})$ are both in the relation $T^{\mathfrak{A}}$, and, moreover, $c_{f(j)+2} = \dots = c_{f(j+1)}$. By transitivity, $(c_{f(j)}, c_{f(j)+2}) \in T^{\mathfrak{A}}$, whence $(c_{f(j)}, c_{f(j+1)}) \in T^{\mathfrak{A}}$. But $c_{f(j)}$ and $c_{f(j+1)}$ are, respectively, the j th and $(j+1)$ th element of \bar{c}^f , and so, by construction, $\mathfrak{A}^+ \models Q[\bar{c}^f]$. Thus $\mathfrak{A}^+ \models \forall \mathbf{x}_k \varphi_T^2$. \dashv

THEOREM 5. *Let $f: [1, k] \rightarrow [1, n]$ be a non-adjacent function. Then the satisfiability and finite satisfiability problems for \mathcal{AF}^f are undecidable.*

PROOF. We reduce from the (finite) satisfiability problem for FO^2 with two transitive relations, known to be undecidable [22, Thm. 3]. Let an FO^2 -sentence φ be given (possibly featuring the two binary predicates T and T' (required to be interpreted as transitive relations), and let φ^* to be its logically equivalent \mathcal{AF} -formula, guaranteed by Theorem 4. Define φ_T and $\varphi_{T'}$ to be the \mathcal{AF}^f -formulas guaranteed by Lemma 18, stating that T and T' are transitive. (The respective predicates Q appearing in these formulas are chosen afresh.) Applying Lemma 18 we see that φ has a (finite) model interpreting T and T' as transitive relations if and only if the \mathcal{AF}^f -formula $\varphi^* \wedge \varphi_T \wedge \varphi_{T'}$ has a (finite) model. \dashv

REFERENCES

- [1] HAJNAL ANDRÉKA, ISTVÁN NÉMETHI, and JOHAN VAN BENTHEM, *Modal Languages and Bounded Fragments of Predicate Logic*, *Journal of Philosophical Logic*, vol. 27 (1998), no. 3, pp. 217–274.
- [2] JAVED A. ASLAM and RONALD L. RIVEST, *Inferring graphs from walks*, *Colt '90: Proceedings of the third annual workshop on computational learning theory* (Mark Fulk and John Case, editors), Morgan Kaufmann, 1990, pp. 359–370.
- [3] FRANZ BAADER, IAN HORROCKS, CARSTEN LUTZ, and ULRIKE SATTLER, *An Introduction to Description Logic*, Cambridge University Press, 2017.
- [4] VINCE BÁRÁNY, BALDER TEN CATE, and LUC SEGOUFIN, *Guarded Negation*, *Journal of the ACM*, vol. 62 (2015), no. 3, pp. 22:1–22:26.
- [5] BARTOSZ BEDNARCZYK, *Exploiting forwardness: Satisfiability and query-entailment in forward guarded fragment*, *Logics in artificial intelligence - 17th european conference, JELIA 2021, virtual event, may 17-20, 2021, proceedings* (Wolfgang Faber, Gerhard Friedrich, Martin Gebser, and Michael Morak, editors), Lecture Notes in Computer Science, vol. 12678, Springer, 2021, pp. 179–193.
- [6] BARTOSZ BEDNARCZYK, WITOLD CHARATONIK, and EMANUEL KIEROŃSKI, *Extending two-variable logic on trees*, *26th EACSL annual conference on computer science logic, CSL 2017, august 20-24, 2017, stockholm, sweden* (Valentin Goranko and Mads Dam, editors), LIPIcs, vol. 82, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, pp. 11:1–11:20.
- [7] BARTOSZ BEDNARCZYK and REIJO JAAKKOLA, *Towards a model theory of ordered logics: Expressivity and interpolation*, *47th international symposium on mathematical foundations of computer science, MFCS 2022, august 22-26, 2022, vienna, austria* (Stefan Szeider, Robert Ganian, and Alexandra Silva, editors), LIPIcs, vol. 241, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 15:1–15:14.
- [8] BARTOSZ BEDNARCZYK, DAUMANTAS KOJELIS, and IAN PRATT-HARTMANN, *On the Limits of Decision: the Adjacent Fragment of First-Order Logic*, *50th international colloquium on automata, languages, and programming (icalp 2023)* (Dagstuhl, Germany) (Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors), Leibniz International Proceedings in Informatics (LIPIcs), vol. 261, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pp. 111:1–111:21.

- [9] BARTOSZ BEDNARCZYK and SEBASTIAN RUDOLPH, *How to Tell Easy from Hard: Complexities of Conjunctive Query Entailment in Extensions of ALC*, *Journal of Artificial Intelligence Research*, vol. 78 (2023).
- [10] MICHAEL BENEDIKT, EGOR KOSTYLEV, and TONY TAN, *Two Variable Logic with Ultimately Periodic Counting*, *SIAM Journal on Computing*, vol. 53 (2024), no. 4, pp. 884–968.
- [11] EGON BÖRGER, ERICH GRÄDEL, and YURI GUREVICH, *The classical decision problem*, Perspectives in Mathematical Logic, Springer, 1997.
- [12] ASHOK K. CHANDRA, DEXTER C. KOZEN, and LARRY J. STOCKMEYER, *Alternation*, *Journal of the ACM*, vol. 28 (1981), no. 1, p. 114–133.
- [13] WITOLD CHARATONIK, EMANUEL KIEROŃSKI, and FILIP MAZOWIECKI, *Decidability of weak logics with deterministic transitive closure*, *Joint meeting of the twenty-third EACSL annual conference on computer science logic (CSL) and the twenty-ninth annual ACM/IEEE symposium on logic in computer science (lics), CSL-LICS '14, vienna, austria, july 14 - 18, 2014* (Thomas A. Henzinger and Dale Miller, editors), ACM, 2014, pp. 29:1–29:10.
- [14] VALENTIN GORANKO and MARTIN OTTO, *Model Theory of Modal Logic*, *Handbook of modal logic* (Patrick Blackburn, J. F. A. K. van Benthem, and Frank Wolter, editors), Studies in logic and practical reasoning, vol. 3, North-Holland, 2007, pp. 249–329.
- [15] ERICH GRÄDEL, *On the Restraining Power of Guards*, *The Journal of Symbolic Logic*, vol. 64 (1999), no. 4, pp. 1719–1742.
- [16] ERICH GRÄDEL, PHOKION G. KOLAITIS, and MOSHE Y. VARDI, *On the decision problem for two-variable first-order logic*, *The Bulletin of Symbolic Logic*, vol. 3 (1997), no. 1, pp. 53–69.
- [17] LEON HENKIN, *Logical systems containing only a finite number of symbols*, Séminaire de mathématiques supérieures, Presses de l'Université de Montréal, 1967.
- [18] ANDREAS HERZIG, *A new decidable fragment of first order logic*, *Abstracts of the 3rd logical biennial summer school and conference in honour of S. C. Kleene* (Varna, Bulgaria), June 1990.
- [19] DAVID HILBERT and WILHELM ACKERMAN, *Grundzüge der theoretischen Logik*, Springer, Berlin, 1928.
- [20] ———, *Principles of mathematical logic*, Chelsea, New York, 1950.
- [21] REIJO JAAKKOLA, *Ordered fragments of first-order logic*, *46th international symposium on mathematical foundations of computer science, MFCS 2021, august 23-27, 2021, tallinn, estonia* (Filippo Bonchi and Simon J. Puglisi, editors), LIPIcs, vol. 202, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, pp. 62:1–62:14.
- [22] EMANUEL KIEROŃSKI, *Results on the Guarded Fragment with Equivalence or Transitive Relations*, *Computer science logic, 19th international workshop, CSL 2005, 14th annual conference of the eacsl, oxford, uk, august 22-25, 2005, proceedings* (C.-H. Luke Ong, editor), Lecture Notes in Computer Science, vol. 3634, Springer, 2005, pp. 309–324.
- [23] EMANUEL KIEROŃSKI and MARTIN OTTO, *Small substructures and decidability issues for first-order logic with two variables*, *The Journal of Symbolic Logic*, vol. 77 (2012), no. 3, pp. 729–765.
- [24] EMANUEL KIEROŃSKI, IAN PRATT-HARTMANN, and LIDIA TENDERA, *Two-variable logics with counting and semantic constraints*, *ACM SIGLOG News*, vol. 5 (2018), no. 3, pp. 22–43.
- [25] EMANUEL KIEROŃSKI and LIDIA TENDERA, *On finite satisfiability of two-variable first-order logic with equivalence relations*, *Proceedings of the 24th annual IEEE symposium on logic in computer science, LICS 2009, 11-14 august 2009, los angeles, ca, USA*, IEEE Computer Society, 2009, pp. 123–132.
- [26] DAUMANTAS KOJELIS, *On Homogenous Models of Fluted Languages*, *Under submission to CSL'25: Proceedings of the 33rd EACSL Annual Conference on Computer Science Logic*, 2025.
- [27] IAN PRATT-HARTMANN, *The two-variable fragment with counting revisited*, *Logic, language, information and computation, 17th international workshop, wollic 2010, brasilia, brazil, july 6-9, 2010. proceedings* (Anuj Dawar and Ruy J. G. B. de Queiroz, editors), Lecture Notes in Computer Science, vol. 6188, Springer, 2010, pp. 42–54.

- [28] ———, *Fluted logic with counting, 48th international colloquium on automata, languages, and programming, ICALP 2021, july 12-16, 2021, glasgow, scotland (virtual conference)* (Nikhil Bansal, Emanuela Merelli, and James Worrell, editors), LIPIcs, vol. 198, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, pp. 141:1–141:17.
- [29] IAN PRATT-HARTMANN, *Fragments of First-Order Logic*, Oxford University Press, 2023.
- [30] IAN PRATT-HARTMANN, *Walking on Words, 35th annual symposium on combinatorial pattern matching, CPM 2024, june 25-27, 2024, fukuoka, japan* (Shunsuke Inenaga and Simon J. Puglisi, editors), LIPIcs, vol. 296, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, pp. 25:1–25:17.
- [31] IAN PRATT-HARTMANN, WIESLAW SZWAST, and LIDIA TENDERA, *The fluted fragment revisited*, *The Journal of Symbolic Logic*, vol. 84 (2019), no. 3, pp. 1020–1048.
- [32] IAN PRATT-HARTMANN and LIDIA TENDERA, *The fluted fragment with transitive relations*, *Annals of Pure and Applied Logic*, vol. 173 (2022), no. 1, p. 103042.
- [33] WILLIAM C. PURDY, *Fluted formulas and the limits of decidability*, *The Journal of Symbolic Logic*, vol. 61 (1996), no. 2, pp. 608–620.
- [34] WILLARD VAN ORMAN QUINE, *On the Limits of Decision*, *Proceedings of the 14th international congress of philosophy*, vol. III, University of Vienna, 1969, pp. 57–62.
- [35] ———, *Algebraic logic and predicate functors*, *The ways of paradox*, Harvard University Press, Cambridge, MA, revised and enlarged ed., 1976, pp. 283–307.
- [36] SYLVAIN SCHMITZ, *Complexity hierarchies beyond elementary*, *ACM Transactions on Computational Logic*, vol. 8 (2016), no. 1.
- [37] THOMAS SCHWENTICK and THOMAS ZEUME, *Two-Variable Logic with Two Order Relations*, *Logical Methods in Computer Science*, vol. Volume 8, Issue 1 (2012).

COMPUTATIONAL LOGIC GROUP, TECHNISCHE UNIVERSITÄT DRESDEN, GERMANY
INSTITUTE OF COMPUTER SCIENCE, UNIVERSITY OF WROCLAW, POLAND
E-mail: bartosz.bednarczyk@cs.uni.wroc.pl

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF MANCHESTER, UK
E-mail: daumantas.kojelis@manchester.ac.uk

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF MANCHESTER, UK
INSTITUTE OF COMPUTER SCIENCE, UNIVERSITY OF OPOLE, POLAND
E-mail: ian.pratt@manchester.ac.uk