CHARACTERIZATIONS OF MONADIC NIP

SAMUEL BRAUNFELD AND MICHAEL C. LASKOWSKI*

Abstract. We give several characterizations of when a complete first-order theory T is monadically NIP, i.e. when expansions of T by arbitrary unary predicates do not have the independence property. The central characterization is a condition on finite satisfiability of types. Other characterizations include decompositions of models, the behavior of indiscernibles, and a forbidden configuration. As an application, we prove non-structure results for hereditary classes of finite substructures of non-monadically NIP models that eliminate quantifiers.

1. Introduction

It is well known that many first order theories whose models are tame can become unwieldy after naming a unary predicate. Arguably the best known example of this is the field $(\mathfrak{C}, +, \cdot)$ of complex numbers. Its theory is uncountably categorical, but after naming a predicate for the real numbers, the expansion becomes unstable. A more extreme example is the theory T of infinite dimensional vector spaces over a finite field, in a relational language. The theory T is totally categorical, but if, in some model V, one names a basis B, then by choosing specified sum sets of basis elements, one can code arbitrary bipartite graphs in expansions of V by unary predicates.

As part of a larger project in [?BS], Baldwin and Shelah undertook a study of this phenomenon. They found that a primary dividing line is whether T *admits coding* i.e., there are three subsets A, B, C of a model of T and a formula $\phi(x, y, z)$ that defines a pairing function $A \times B \to C$. If one can find such a configuration in a model M of T, some monadic expansions of M are wild. The primary focus in [?BS] was *monadically stable* theories, i.e. theories that remain stable after arbitrary expansions by unary predicates. Clearly, the two theories described above are stable, but not monadically stable. They offered a characterization of monadically stable theories within the stable theories via a condition on the behavior of non-forking. This allowed them to prove that monadic stability yields a dividing line within stable theories: models of monadically stable theories are well-structured and admit a nice decomposition into trees of submodels, while if a theory is stable but not monadically stable then it encodes arbitrary bipartite graphs in a unary expansion, and so is not even monadically NIP.

1991 *Mathematics Subject Classification.* 03C45.

*Partially supported by NSF grant DMS-1855789.

A theory T is NIP if it does not have the independence property, and is monadically NIP if every expansion of a model of T by unary predicates is also NIP. The behavior of NIP theories has been extensively studied, see e.g., [?Pierre]. Soon after [?BS], Shelah further studied monadically NIP theories in [?Hanf], where he showed they satisfy a condition on the behavior of finite satisfiable types paralleling the condition on the behavior of non-forking in monadically stable theories. He was then able to use this to produce a linear decomposition of models of monadically NIP theories, akin to a single step of the tree decomposition in monadically stable theories.

We dub Shelah's condition on the behavior of finite satisfiability the *f.s. dichotomy*, and we consider it to be the fundamental property expressible in the original language L describing the dichotomous behavior outlined above. We show the f.s. dichotomy characterizes monadically NIP theories and provide several other characterizations, including admitting a linear decomposition in the style of Shelah, a forbidden configuration, and conditions on the behavior of indiscernible sequences after adding parameters. Definitions for the following theorem may be found in Definitions 3.9, 3.1, 3.4, and 3.8. Of note is that all but the first two conditions refer to the theory T itself, rather than unary expansions.

Theorem 1.1. *The following are equivalent for a complete theory T with an infinite model.*

(1) *T is monadically NIP.*
(2) *No monadic expansion of T admits coding.*
(3) *T does not admit coding on tuples.*
(4) *T has the f.s. dichotomy.*
(5) *For all $M^* \models T$ and $M, N \preceq M^*$, every partial M-f.s. decomposition of N extends to an (irreducible) M-f.s. decomposition of N.*
(6) *T is dp-minimal and has indiscernible-triviality.*

We believe that monadic NIP (or perhaps a quantifier-free version) is an important dividing line in the combinatorics of hereditary classes, and provides a general setting for the sort of decomposition arguments common in structural graph theory. For example, see the recent work on bounded twin-width in the ordered binary case, where it coincides with monadic NIP [?ST, ?TW4]. Here, we mention the following conjecture, adding monadic NIP to a question of Macpherson [?MacHom, Question 2.2.7].

Conjecture 1. *The following are equivalent for a countable homogeneous ω-categorical relational structure M.*

(1) *M is monadically NIP.*
(2) *The (unlabeled) growth rate of $Age(M)$ is at most exponential.*
(3) *$Age(M)$ is well-quasi-ordered under embeddability, i.e. it has no infinite antichain.*

From Theorem 1.1, we see that if T is not monadically NIP then it admits coding on tuples. This allows us to prove the following non-structure theorem in §5 (with Definition 5.1 defining the relevant terms), in particular confirming (2) ⇒ (1) and a weak form of (3) ⇒ (1) from the conjecture, although without any assumption of ω-categoricity.

Theorem 1.2. *Suppose T is a complete theory with quantifier elimination in a relational language with finitely many constants. If T is not monadically NIP, then $Age(T)$ has growth rate asymptotically greater than $(n/k)!$ for some $k \in \omega$ and is not 4-wqo.*

We also show the following, partially explaining the importance of monadic model-theoretic properties for the study of hereditary classes.

Theorem 1.3. *Suppose T is a complete theory with quantifier elimination in a relational language with finitely many constants. Then $Age(T)$ is NIP if and only if T is monadically NIP, and $Age(T)$ is stable if and only if T is monadically stable.*

In Section 2, we review basic facts about finite satisfiability, and introduce *M-f.s. sequences*, which are closely related to, but more general than Morley sequences. The results of this section apply to an arbitrary theory, and so may well be of interest beyond monadic NIP. Section 3 introduces the f.s. dichotomy and proves the equivalence of (3)-(6) from Theorem 1.1. Much of these two sections is an elaboration on the terse presentation of [?Hanf], although there are new definitions and results, particularly in Section 3.2, which deals with the behavior of indiscernibles in monadically NIP theories. In Section 4 we finish proving the main theorem by giving a type-counting argument that the f.s. dichotomy implies monadic NIP, and by showing that if T admits coding on tuples then it admits coding in a unary expansion. In Section 5, we prove Theorems 1.2 and 1.3.

We are grateful to Pierre Simon, with whom we have had numerous insightful discussions about this material. In particular, the relationship between monadic NIP and indiscernible-triviality was suggested to us by him.

1.1. **Notation.** Throughout this paper, we work in $\mathfrak{C}$, a large, sufficiently saturated, sufficiently homogeneous model of a complete theory T. We routinely consider $\text{tp}(A/B)$ when A is an infinite set. To make this notion precise, we (silently) fix an enumeration $\bar{a}$ of A (of ordinal order type) and an enumeration $\bar{x}$ with $\lg(\bar{x}) = \lg(\bar{a})$. Then $\text{tp}(A/B) = \{\theta(\bar{x}', \bar{b}) : \mathfrak{C} \models \theta(\bar{a}', \bar{c})$ for all subsequences $\bar{x}' \subseteq \bar{x}$ and $\bar{a}'$ is the corresponding subsequence of $\bar{a}\}$.

2. M-f.s. sequences

Forking independence and Morley sequences are fundamental tools in the analysis of monadically stable theories in [?BS]. These are less well-behaved outside the stable setting, but in any theory we may view '$\text{tp}(A/MB)$ is finitely satisfiable in M' as a statement that A is (asymmetrically) independent from B over M. Following [?Hanf], we will use finite satisfiability in place of non-forking, and indiscernible M-.f.s. sequences in place of Morley sequences. Throughout Section 2, we make no assumptions about the complexity of $Th(\mathfrak{C})$.

2.1. **Preliminary facts about M-f.s. sequences.** For the whole of this section, fix a small $M \preceq \mathfrak{C}$ (typically, $|M| = |T|$).

Definition 2.1. Suppose $B \supseteq M$. Then for any A (possibly infinite) we say $\mathrm{tp}(A/B)$ is *finitely satisfied in M* if, for all $\theta(\bar{y}, \bar{b}) \in \mathrm{tp}(A/B)$, there is $\bar{m} \in M^{\lg(\bar{y})}$ such that $\mathfrak{C} \models \theta(\bar{m}, \bar{b})$.

One way of producing finitely satisfiable types in M comes from *average types.*

Definition 2.2. Suppose $\bar{x}$ is a possibly infinite tuple. For any ultrafilter $\mathcal{U}$ on $M^{\lg(\bar{x})}$ and any $B \supseteq M$,

$$Av(\mathcal{U}, B) = \{\, \phi(\bar{x}, \bar{b}) : \{\, \bar{m} \in M^{\lg(\bar{x})} : \mathfrak{C} \models \phi(\bar{m}, \bar{b}) \,\} \in \mathcal{U} \,\}$$

It is easily checked that $Av(\mathcal{U}, B)$ is a complete type over B that is finitely satisfied in M. We record a few basic facts about types that are finitely satisfied in M. Proofs can be found in either Section VII.4 of [?Shc] or in [?Pierre].

Fact 2.3. *Let M be any model.*

(1) *For any set $B \supseteq M$ and any $p(\bar{x}) \in S(B)$ ($\bar{x}$ may be an infinite tuple), p is finitely satisfied in M if and only if $p = Av(\mathcal{U}, B)$ for some ultrafilter $\mathcal{U}$ on $M^{\lg(\bar{x})}$.*
(2) *Suppose $\Gamma(\bar{x}, B)$ is any set of formulas, closed under finite conjunctions, and each of which is realized in M. Then there is a complete type $p \in S(B)$ extending Γ that is finitely satisfied in M.*
(3) *(Non-splitting) If $p \in S(B)$ is finitely satisfied in M, then p does not split over M, i.e., if $\bar{b}, \bar{b}' \subseteq B$ and $\mathrm{tp}(\bar{b}/M) = \mathrm{tp}(\bar{b}'/M)$, then for any $\phi(\bar{x}, \bar{y})$, $\phi(\bar{x}, \bar{b}) \in p$ if and only if $\phi(\bar{x}, \bar{b}') \in p$.*
(4) *(Transitivity) If $\mathrm{tp}(B/C)$ and $\mathrm{tp}(A/BC)$ are both finitely satisfied in M, then so is $\mathrm{tp}(AB/C)$.*

Definition 2.4 (M-f.s. sequence)**.** With M fixed as above, let $(I, \leq)$ be any linearly ordered index set.

- Suppose $\langle A_i : i \in I \rangle$ is any sequence of sets, indexed by $(I, \leq)$. For $J \subseteq I$, A_J denotes $\bigcup_{j \in J} A_j$, and for $i^* \in I$, $A_{<i^*}$ denotes $\bigcup_{i < i^*} A_i$. $A_{\leq i^*}$ and $A_{>i^*}$ are defined analogously.
- For $C \supseteq M$, an *M-f.s. sequence over C*, is a sequence of sets $\langle A_i : i \in I \rangle$ such that $\mathrm{tp}(A_i / A_{<i} C)$ is finitely satisfied in M for every $i \in I$. When $C = M$ we simply say $\langle A_i : i \in I \rangle$ is an M-f.s. sequence.

Note that for any $C \supseteq M$, $\langle A_i : i \in I \rangle$ is an M-f.s. sequence over C if and only if the concatenation $\langle C \rangle \frown \langle A_i : i \in I \rangle$ is an M-f.s. sequence.

We note two useful operations on M-f.s. sequences over C, 'Shrinking' and 'Condensation'.

Definition 2.5. Suppose $C \supseteq M$ and $\langle A_i : i \in I\rangle$ is an M-f.s. sequence over C.

(1) 'Shrinking:' For every $J \subseteq I$, for all $A'_j \subseteq A_j$, and for all C' with $C \supseteq C' \supseteq M$, we say $\langle A'_j : j \in J\rangle$ as a sequence over C' is obtained by *shrinking* from $\langle A_i : i \in I\rangle$ as a sequence over C.
(2) Condensation:' Suppose $\pi : I \to J$ is a condensation, i.e., a surjective map with each $\pi^{-1}(j)$ a convex subset of I. For each $j \in J$, let $A^*_j := \bigcup \{ A_i : i \in \pi^{-1}(j) \}$. We say $\langle A^*_j : j \in J\rangle$ as a sequence over C is obtained by *condensation* from $\langle A_i : i \in I\rangle$ as a sequence over C.

In particular, removing a set of A_i's from the sequence is an instance of Shrinking.

Lemma 2.6. *Suppose $C \supseteq M$ and $\langle A_i : i \in I\rangle$ is an M-f.s. sequence over C. Then Shrinking and Condensation both preserve being an M-f.s. sequence over C. In particular, for any partition $I = J \sqcup K$ into convex pieces, the two-element sequence $\langle A_J, A_K\rangle$ is an M-f.s. sequence over C.*

Proof. The statement is immediate for Shrinking, and for Condensation follows by transitivity in Fact 2.3. The last sentence is a special case of Condensation, as the partition defines a condensation $\pi : I \to \{ 0, 1 \}$ with $\pi^{-1}(0) = J$. □

Definition 2.7. If $\langle A_i : i \in I\rangle$ is an M-f.s. sequence over C, call $\langle B_j : j \in J\rangle$ a *simple extension, resp. blow-up* if $\langle A_i : i \in I\rangle$ is attained from it by Shrinking, resp. by Condensation. $\langle D_k : k \in K\rangle$ is an *extension* of $\langle A_i : i \in I\rangle$ if it is a blow-up of a simple extension of $\langle A_i : i \in I\rangle$ over C.

Here is one general result, whose verification is just bookkeeping.

Lemma 2.8. *Suppose $\langle A_i : i \in I\rangle$ is an M-f.s. sequence, $i^* \in I$, $J \cap I = \emptyset$, and $\langle A'_j : j \in J\rangle/MA_{<i^*}$ is an M-f.s. sequence over $MA_{<i^*}$ with $\bigcup \{ A'_j : j \in J \} = A_{i^*}$. Then the blow-up*

$$\langle A_i : i < i^*\rangle \frown \langle A'_j : j \in J\rangle \frown \langle A_i : i > i^*\rangle$$

is also an M-f.s. sequence.

The next lemma is not used later, but shows that if $M \preceq N$, then decomposing N as an M-f.s. sequence gives a chain of elementary substructures approximating N.

Lemma 2.9. *Suppose $M \preceq N$ and $\langle A_i : i \in I\rangle$ is any M-f.s. sequence with $MA_I = N$. Then, for every initial segment $I_0 \subseteq I$ (regardless of whether or not I_0 has a maximum) MA_{I_0} is an elementary substructure of N.*

Proof. We apply the Tarski-Vaught criterion. Choose a formula $\phi(x, \bar{a}, \bar{m})$ with $\bar{a}$ from A_{I_0} and $\bar{m}$ from M such that $N \models \exists x \phi(x, \bar{a}, \bar{m})$. If some $c \in N \setminus MA_{I_0}$ realizes $\phi(x, \bar{a}, \bar{m})$, then as $\mathrm{tp}(c/MA_{I_0})$ is finitely satisfied in

M, there is also a solution in M. Otherwise, if there is a solution in MA_{I_0}, there is nothing to check. $\square$

The following argument is contained in the proof of [?Hanf, Part I Lemma 2.6], but the statement here is slightly more general. (The paper [?Hanf] is divided into Part I and Part II, with overlapping numbering schemes.)

Proposition 2.10 (Extending the base)**.** *Suppose $C \supseteq M$ and $\langle A_i : i \in I\rangle$ is an M-f.s. sequence over C. For every $D \supseteq C$, there is D' with $\mathrm{tp}(D'/C) = \mathrm{tp}(D/C)$ and $\langle A_i : i \in I\rangle$ is an M-f.s. sequence over D'.*

Proof. As notation, choose disjoint sets $\{ \bar{x}_i : i \in I \}$ of variables, with $\lg(\bar{x}_i) = \lg(A_i)$ for each $i \in I$. For each $i \in I$, choose an ultrafilter $\mathcal{U}_i$ on $M^{\lg(A_i)}$ such that $\mathrm{tp}(A_i/A_{<i}C) = Av(\mathcal{U}_i, A_{<i}C)$.

For a finite, non-empty $t = \{ i_1 < i_2 < \cdots < i_n \} \subseteq I$, let $\bar{x}_t = \bar{x}_{i_1} \dots \bar{x}_{i_n}$. We will recursively define complete types $w_t(\bar{x}_t) \in S_{\bar{x}_t}(D)$ as follows:

- For $t = \{ i^* \}$ a singleton, let $w_t(\bar{x}_t) := Av(\mathcal{U}_{i^*}, D)$.
- For $|t| > 1$, letting $i^* = \max(t)$ and $s = t \setminus \{ i^* \}$,

$$w_t(\bar{x}_t) := w_s(\bar{x}_s) \cup Av(\mathcal{U}_i, D\bar{x}_s)$$

That is, $\bar{a}'_t$ realizes w_t if and only if $\bar{a}'_s$ realizes w_s and, for every $\theta(\bar{x}_{i^*}, \bar{d}, \bar{a}'_s)$, $\theta(\bar{a}'_{i^*}, \bar{d}, \bar{a}'_s)$ holds if and only if $\{ \bar{m} \in M^{\lg(A_{i^*})} : \mathfrak{C} \models \theta(\bar{m}, \bar{d}, \bar{a}'_s) \} \in \mathcal{U}_{i^*}$.

It is easily checked that each $w_t(\bar{x}_t)$ is a complete type over D and, arguing by induction on $|t|$, whenever $t' \subseteq t$, $w_{t'}$ is the restriction of w_t to $\bar{x}_{t'}$. Thus, by compactness, $w^* := \bigcup \{ w_t(\bar{x}_t) : t \subseteq I \text{ non-empty, finite} \}$ is consistent, and in fact, is a complete type over D. Choose any realization $\langle A'_i : i \in I\rangle$ of w^*. Then, for each $i \in I$, $\mathrm{tp}(A'_i/DA'_{<i}) = Av(\mathcal{U}_i, DA'_{<i})$. Since $D \supseteq C$ and $\mathrm{tp}(A_i/CA_{<i}) = Av(\mathcal{U}_i, CA_{<i})$, it follows that $\mathrm{tp}(\langle A'_i : i \in I\rangle/C) = \mathrm{tp}(\langle A_i : i \in I\rangle/C)$. Thus, it suffices to choose any D' satisfying $\langle A'_i : i \in I\rangle D \equiv_C \langle A_i : i \in I\rangle D'$. $\square$

2.2. $C \supseteq M$ full for non-splitting.

Definition 2.11. We call $C \supseteq M$ *full* (for non-splitting over M) if, for every n, every $p \in S_n(M)$ is realized in C.

The relevance of fullness is that, whenever $C \supseteq M$ is full, every complete type $q \in S(C)$ has a unique extension to any set $D \supseteq C$ that does not split over M. Keeping in mind finite satisfiability as an analogue of non-forking, the next lemma says that 'types over C that are finitely satisfied in M are stationary.'

Lemma 2.12 ([?Hanf, Part I Lemma 1.5])**.** *Suppose $C \supseteq M$ is full and $p \in S(C)$ is finitely satisfied in M. Then for any set $D \supseteq C$, there is a unique $q \in S(D)$ extending p that remains finitely satisfied over M.*

Proof. In fact, we can easily describe q. A formula $\theta(\bar{x}, \bar{d}) \in q$ if and only if $\theta(\bar{x}, \bar{c}) \in p$ for some (equivalently, for every) $\bar{c}$ from C with $\mathrm{tp}(\bar{d}/M) = \mathrm{tp}(\bar{c}/M)$. The fact that q is well-defined is because, being finitely satisfied in M, p does not split over M. $\square$

Lemma 2.13 ([?Hanf, Part I Observation 1.6])**.** *Suppose $C \supseteq M$ is full and $\langle A, B\rangle/C$ is an M-f.s. sequence over C. Partition B as B_1B_2 (not necessarily convex). Then $\langle A, B_1, B_2\rangle/C$ is an M-f.s. sequence over C if and only if $\langle B_1, B_2\rangle/C$ is an M-f.s. sequence over C.*

Proof. Left to right is obvious. For the converse, we need to show that $\mathrm{tp}(B_2/B_1AC)$ is finitely satisfied in M. To begin, by Proposition 2.10, choose $B_2' \equiv_{B_1C} B_2$ with $\mathrm{tp}(B_2'/B_1AC)$ finitely satisfied in M. Note that

$$B_1B_2' \equiv_C B_1B_2$$

Also, since $\langle A, B\rangle/C$ is an M-f.s. sequence over C, we have both $\langle A, B_1B_2\rangle/C$ and (by Shrinking) $\langle A, B_1\rangle/C$ are M-f.s. sequences over C. By transitivity, the last statement, coupled with $\mathrm{tp}(B_2'/B_1AC)$ finitely satisfied in M, implies $\mathrm{tp}(B_1B_2'/AC)$ is finitely satisfied in M. Thus, by Lemma 2.12,

$$B_1B_2' \equiv_{AC} B_1B_2$$

As $\mathrm{tp}(B_2'/B_1AC)$ finitely satisfied in M, so is $\mathrm{tp}(B_2/B_1AC)$. □

Lemma 2.14. *Suppose $C \supseteq M$ is full and $\langle A, B\rangle/C$ is an M-f.s. sequence over C. Choose any $\bar{a}_1, \bar{a}_2$ from A and $\bar{b}_1, \bar{b}_2$ from B with $\mathrm{tp}(\bar{a}_1/C) = \mathrm{tp}(\bar{a}_2/C)$ and $\mathrm{tp}(\bar{b}_1/C) = \mathrm{tp}(\bar{b}_2/C)$. Then $\mathrm{tp}(\bar{a}_1\bar{b}_1/C) = \mathrm{tp}(\bar{a}_2\bar{b}_2/C)$.*

Proof. Let $p = \mathrm{tp}(\bar{a}_1/C)$. As $\mathrm{tp}(\bar{a}_1/C) = \mathrm{tp}(\bar{a}_2/C)$, the map $f : C\bar{a}_1 \to C\bar{a}_2$ fixing C pointwise with $f(\bar{a}_1) = \bar{a}_2$ is elementary. To prove the Lemma, it suffices to show that $\bar{b}_2$ realizes $f(p)$.

To see this, let $\bar{b}^*$ be any realization of $f(p)$ (anywhere in $\mathfrak{C}$). Then $\mathrm{tp}(\bar{a}_1\bar{b}_1/C) = \mathrm{tp}(\bar{a}_2\bar{b}^*/C)$. From this it follows that $\mathrm{tp}(\bar{b}^*/C) = \mathrm{tp}(\bar{b}_1/C) = \mathrm{tp}(\bar{b}_2/C)$, with the second equality by hypothesis. But also:

(1) $\mathrm{tp}(\bar{b}^*/C\bar{a}_2)$ is finitely satisfied in M since $\mathrm{tp}(\bar{a}_1\bar{b}_1/C) = \mathrm{tp}(\bar{a}_2\bar{b}^*/C)$ and $\mathrm{tp}(B/AC)$ is finitely satisfied in M; and
(2) $\mathrm{tp}(\bar{b}_2/C\bar{a}_2)$ is finitely satisfied in M since $\mathrm{tp}(B/AC)$ is finitely satisfied in M.

Applying Lemma 2.12 to the last three statements implies $\mathrm{tp}(\bar{b}_2/\bar{a}_2C) = \mathrm{tp}(\bar{b}^*/\bar{a}_2C)$, i.e., $\bar{b}_2$ realizes $f(p)$. □

We glean two results from Lemma 2.14. The first bounds the number of types realized in an M-f.s. sequence, independent of either $|I|$ or $|A_i|$.

Lemma 2.15. *For any model M, for any M-f.s. sequence $\langle A_i : i \in I\rangle$, and for every $i^* \in I$, $k \in \omega$, the number of complete k-types over $A_{<i^*}M$ realized in $A_{\geq i^*}$ is at most $\beth_2(|M|)$.*

Proof. Because of Condensation, it suffices to prove that for any model M and for any M-f.s. sequence $\langle A, B\rangle$, at most $\beth_2(|M|)$ complete k-types are realized in B. To see this, choose a full $C_0 \supseteq M$ with $|C_0| \leq 2^{|M|}$. By Proposition 2.10, choose $C \supseteq M$ with $\mathrm{tp}(C/M) = \mathrm{tp}(C_0/M)$ and $\langle A, B\rangle$ an M-f.s. sequence over C. Choose any $\bar{b}, \bar{b}' \in B^k$ with $\mathrm{tp}(\bar{b}/C) = \mathrm{tp}(\bar{b}'/C)$. As both $\mathrm{tp}(\bar{b}/AC)$ and $\mathrm{tp}(\bar{b}'/AC)$ are finitely satisfied in M, it follows from

Lemma 2.12 that $\operatorname{tp}(\bar{b}/AC) = \operatorname{tp}(\bar{b}'/AC)$. As there are at most $\beth_2(|M|)$ complete k-types over C, this suffices. $\square$

The second is a refinement of the type structure of an M-f.s. sequence over a full $C \supseteq M$.

Definition 2.16. An M-f.s. sequence $\langle A_i : i \in I\rangle/C$ is an *order-congruence over C* if, for all $i^* \in I$, for all $i^* \leq i_1 < i_2 \cdots < i_n$, $i^* \leq j_1 < j_2 < \dots j_n$ from I, and for all $\bar{a}_k \in A_{i_k}, \bar{b}_k \in A_{j_k}$ satisfying $\operatorname{tp}(\bar{a}_k/C) = \operatorname{tp}(\bar{b}_k/C)$ for $k = 1, \dots, n$, we have

$$\operatorname{tp}(\bar{a}_1, \dots, \bar{a}_n/CA_{<i^*}) = \operatorname{tp}(\bar{b}_1, \dots, \bar{b}_n/CA_{<i^*})$$

The following is essentially part of the statement of [?Hanf, Part I Lemma 2.6].

Proposition 2.17. *For every model M, every M-f.s. sequence $\langle A_i : i \in I\rangle$ over any full $C \supseteq M$ is an order-congruence over C.*

Proof. Fix i^*, $i^* \leq i_1 < \dots, i_n$, $i^* \leq j_1 < \cdots < j_n$, $\bar{a}_1, \dots, \bar{a}_n$, and $\bar{b}_1, \dots, \bar{b}_n$ as in the hypotheses. By shrinking, for each $1 \leq k < n$, both $\operatorname{tp}(\bar{a}_{k+1}/C\bar{a}_1, \dots, \bar{a}_k)$ and $\operatorname{tp}(\bar{b}_{k+1}/C\bar{b}_1, \dots, \bar{b}_k)$ are finitely satisfied in M. As $C \supseteq M$ is full, by iterating Lemma 2.14 $(n-1)$ times, we have $\operatorname{tp}(\bar{a}_1, \dots, \bar{a}_n/C) = \operatorname{tp}(\bar{b}_1, \dots, \bar{b}_n/C)$. Also, both $\operatorname{tp}(\bar{a}_1, \dots, \bar{a}_n/CA_{<i^*})$ and $\operatorname{tp}(\bar{b}_1, \dots, \bar{b}_n/CA_{<i^*})$ are finitely satisfied in M, so $\operatorname{tp}(\bar{a}_1, \dots, \bar{a}_n/CA_{<i^*}) = \operatorname{tp}(\bar{b}_1, \dots, \bar{b}_n/CA_{<i^*})$ by Lemma 2.12. $\square$

2.3. **M-f.s. sequences and indiscernibles.** In this subsection, we explore the relation between M-f.s. sequences and indiscernibles. An M-f.s. sequence need not be indiscernible (for example, the tuples can realize different types), but when it is, it gives a special case of a Morley sequence in the sense of [?Pierre].

We first show indiscernible sequences can always be viewed as M-f.s. sequences over some model M.

Lemma 2.18 (extending [?Hanf, Part I Lemma 4.1])**.** *Suppose $(I, \leq)$ is infinite and $\mathcal{I} = \langle \bar{a}_i : i \in I\rangle$ is indiscernible over $\emptyset$. (For simplicity, assume $\lg(\bar{a}_i)$ is finite). Then there is a model M such that $\langle \bar{a}_i : i \in I\rangle$ is both indiscernible over M and is an M-f.s. sequence.*

Furthermore, if there is some B such that $\mathcal{I}$ is indiscernible over each $b \in B$, then M may be chosen so that $\mathcal{I}$ additionally remains indiscernible over Mb.

Proof. Expand the language to have built-in Skolem functions while keeping $\mathcal{I}$ indiscernible, and end-extend $\mathcal{I}$ to an indiscernible sequence of order-type $I + \omega^*$. (For the 'Furthermore' sentence, note this can still be done so the result is indiscernible over each $b \in B$.) Let I^* be the new elements added, and let M be reduct of the Skolem hull of I^* to the original language. (If $I + I^*$ were indiscernible over b, then I is indiscernible over Mb.) $\square$

Armed with this Lemma, we characterize when an infinite $\mathcal{I} = \langle \bar{a}_i : i \in I \rangle$ is both an M-f.s sequence and is indiscernible over M. (A paradigm of an indiscernible sequence over M that is not an M-f.s. sequence is where M is an equivalence relation with infinitely many, infinite classes and $\langle a_i : i \in \omega \rangle$ is a sequence from some E-class not represented in M.)

Lemma 2.19. *An infinite sequence $\mathcal{I} = \langle \bar{a}_i : i \in I \rangle$ of n-tuples is both indiscernible over M and an M-f.s. sequence if and only if there is an ultrafilter $\mathcal{U}$ on M^n such that $\mathrm{tp}(\bar{a}_i/MA_{<i}) = Av(\mathcal{U}, MA_{<i})$ for every $i \in I$.*

Proof. Right to left is clear, so assume $\mathcal{I}$ is both indiscernible over M and an M-f.s. sequence. As $(I, \leq)$ is infinite, it contains either an ascending or descending ω-chain. For definiteness, choose $J \subseteq I$ of order type ω. To ease notation, we write $\bar{a}_k$ in place of $\bar{a}_{j_k}$. For each $k \in \omega$ and each formula $\phi(\bar{x}, \bar{b}) \in \mathrm{tp}(\bar{a}_k/MA_{<k})$, let $S^k_{\phi(\bar{x},\bar{b})} = \{ \bar{m} \in M^n : \mathfrak{C} \models \phi(\bar{a}_k, \bar{b}) \}$. As $\langle \bar{a}_k : k \in \omega \rangle$ is indiscernible over M, $S^k_{\phi(\bar{x},\bar{b})} = S^\ell_{\phi(\bar{x},\bar{b})}$ for all $\ell \geq k$, and because it is an M-f.s. sequence, $\bigcup \{ S^k_{\phi(\bar{x},\bar{b})} : k \in \omega, \phi(\bar{x}, \bar{b}) \in \mathrm{tp}(\bar{a}_k/MA_{<k}) \}$ has the finite intersection property. Choose any ultrafilter $\mathcal{U}$ on M^n containing every $S^k_{\phi(x,\bar{b})}$. Thus, for any $k \leq \ell < \omega$ and $\phi(\bar{x}, \bar{b})$ with $\bar{b} \subseteq A_{<k}$,

$$\phi(\bar{x}, \bar{b}) \in \mathrm{tp}(\bar{a}_\ell/MA_{<\ell}) \quad \Leftrightarrow \quad S_{\phi(\bar{x},\bar{b})} \in \mathcal{U} \quad \Leftrightarrow \quad \phi(\bar{x}, \bar{b}) \in Av(\bar{a}_\ell/MA_{<\ell})$$

Finally, as $J \subseteq I$ and $\mathcal{I}$ is indiscernible over M, an easy induction on $\lg(\bar{b})$ gives the result. □

Using Lemma 2.19, we obtain a strengthening of Lemma 2.18. The lemma below can be proved by modifying the proof of Lemma 2.10, but the argument here is fundamental enough to bear repeating.

Lemma 2.20. *If an infinite $\mathcal{I} = \langle \bar{a}_i : i \in I \rangle$ is both indiscernible over M and an M-f.s. sequence, then for any $C \supseteq M$, there is $C' \supseteq M$ such that $\mathrm{tp}(C'/M) = \mathrm{tp}(C/M)$ and $\mathcal{I}$ is both indiscernible over C' and an M-f.s. sequence over C'. Thus, if $\mathcal{I}$ is an infinite, indiscernible sequence over $\emptyset$, then there is a model M and a full $C \supseteq M$ such that $\mathcal{I}$ is both indiscernible over C and an M-f.s. sequence over C.*

Proof. For the first sentence, given $\mathcal{I}$, M and C, choose an ultrafilter $\mathcal{U}$ as in Lemma 2.19. A routine compactness argument shows that we can find a sequence $\langle \bar{a}^*_i : i \in I \rangle$ such that $\mathrm{tp}(\bar{a}^*_i/CA_{<i}) = Av(\mathcal{U}, CA^*_{<i})$ for every $i \in I$. As we also have $\mathrm{tp}(\bar{a}_i : MA_{<i}) = Av(\mathcal{U}, MA_{<i})$, an easy induction shows that $\mathrm{tp}(I/M) = \mathrm{tp}(I^*/M)$. Now any C' satisfying $\mathrm{tp}(IC'/M) = \mathrm{tp}(I^*C/M)$ suffices.

For the second sentence, given $\mathcal{I}$, apply Lemma 2.18 to get an M for which $\mathcal{I}$ is both indiscernible over M and an M-f.s. sequence, and choose any full $C \supseteq M$. Then apply the first sentence to $\mathcal{I}$, M, and C. □

Next, we recall the following characterization of indiscernibility. The relevant concepts first appeared in the proof of Theorem 4.6 of [?Morley] and a full proof appears in [?Shc, Lemma I, 2.5].

Lemma 2.21. *Suppose $\langle A_i : i \in I\rangle$ is any sequence of sets indexed by a linear order $(I, \leq)$ and let B be any set. For each $i \in I$, fix a (possibly infinite) enumeration $\bar{a}_i$ of A_i and let $p_i(\bar{x}) = \mathrm{tp}(\bar{a}_i/BA_{<i})$. Then $\langle \bar{a}_i : i \in I\rangle$ is indiscernible over B if and only if*

(1) For all $i \leq j$, $\bar{a}_j$ realizes p_i; and
(2) Each p_i does not split over B.

By contrast, if $C \supseteq M$ and $\langle A_i : i \in I\rangle/C$ is an M-f.s. sequence over C, then (2) is satisfied, but (1) may fail. In the case where $C \supseteq M$ is full, (1) reduces to a question about types over C.

Lemma 2.22. *Suppose $C \supseteq M$ is full and $\langle A_i : i \in I\rangle$ is an M-f.s. sequence over C. Then $\langle A_i : i \in I\rangle$ is indiscernible over C if and only if $\mathrm{tp}(A_i/C) = \mathrm{tp}(A_j/C)$ for all $i, j \in I$.*

Proof. Left to right is clear. For the converse, fix $i < j$. By Lemma 2.21, it suffices to show A_j realizes p_i. But this is clear, as both $\mathrm{tp}(A_i/A_{<i}C)$ and $\mathrm{tp}(A_j/A_{<i}C)$ are finitely satisfied in M and $\mathrm{tp}(A_i/C) = \mathrm{tp}(A_j/C)$. Since $C \supseteq M$ is full, $\mathrm{tp}(A_j/A_{<i}C) = \mathrm{tp}(A_i/A_{<i}C) = p_i$ by Lemma 2.12. □

In terms of existence of such sequences, we have the following.

Lemma 2.23. *Suppose $C \supseteq M$ is full and $p(\bar{x}) \in S(C)$ is finitely satisfied in M. Then for every $(I, \leq)$ there is an M-f.s. sequence $\langle \bar{a}_i : i \in I\rangle$ over C of realizations of p, hence is also indiscernible over C.*

Proof. By compactness it suffices to prove this for $(I, \leq) = (\omega, \leq)$. By Fact 2.3(1), choose an ultrafilter $\mathcal{U}$ on $M^{\mathrm{lg}(\bar{x})}$ and recursively let $\bar{a}_i$ be a realization of $Av(\mathcal{U}, C\bar{a}_{<i})$. It is easily checked that $\langle \bar{a}_i : i \in \omega\rangle/C$ is an M-f.s. sequence over C with $\mathrm{tp}(\bar{a}_i/C) = p$ for each i. As $C \supseteq M$ is full, it is also indiscernible over C by Lemma 2.22. □

3. The f.s. dichotomy

We begin this section with the central dividing line of this paper. Although unnamed, the concept appears in Lemma II 2.3 of [?Hanf].

Definition 3.1 (f.s. dichotomy)**.** T has the *f.s. dichotomy* if, for all models M, all finite tuples $\bar{a}, \bar{b} \in \mathfrak{C}$, if $\mathrm{tp}(\bar{b}/M\bar{a})$ is finitely satisfied in M, then for any singleton c, either $\mathrm{tp}(\bar{b}/M\bar{a}c)$ or $\mathrm{tp}(\bar{b}c/M\bar{a})$ is finitely satisfied in M.

It would be equivalent to replace $\bar{a}$, $\bar{b}$ by sets $A, B \subset \mathfrak{C}$ in the definition above, and this form will often be used. Much of the utility of the f.s. dichotomy is via the following extension lemma.

Lemma 3.2 ([?Hanf, Part I Claim 2.4])**.** *Suppose T has the f.s. dichotomy and $\langle A_i : i \in I\rangle$ is any M-f.s. sequence. Then for every $c \in \mathfrak{C}$, there is a simple extension $\langle A'_j : j \in J\rangle$ of $\langle A_i : i \in I\rangle$ that includes c that is also an M-f.s. sequence. Moreover, if $(I, \leq)$ is a well-ordering with a maximum element, we may take $J = I$.*

Proof. Fix any M-f.s. sequence $\langle A_i : i \in I\rangle$ and choose any singleton $c \in \mathfrak{C}$. Let $I_0 \subseteq I$ be the maximal initial segment of I such that $\mathrm{tp}(c/A_{I_0}M)$ is finitely satisfied in M. Note that I_0 could be empty or all of I. If the minimum element of $(I \setminus I_0)$ exists, name it i^* and take $J = I$; otherwise, let $J = I \cup \{i^*\}$, where i^* is a new element realizing the cut $(I_0, I \backslash I_0)$ and put $A_{i^*} = \emptyset$.

Let $A'_{i^*} := A_{i^*} \cup \{ c \}$ and $A'_i := A_i$ for all $i \neq i^*$. We claim that $\langle A'_i : i \in J\rangle$ is an M-f.s. sequence. To see this, note that $A'_{<i^*} = A_{<i^*} = A_{I_0}$ while $A'_{<j}$ properly extends A_{I_0} for any $j > i^*$. Thus, $\mathrm{tp}(c/A_{<i^*}M)$ is finitely satisfied in M, but $\mathrm{tp}(c/A_{<j}M)$ is not for every $j > i^*$.

We first show that $\mathrm{tp}(A_{i^*}c/A_{<i^*}M)$ is finitely satisfied in M. This is immediate if $A_{i^*} = \emptyset$. If not and $A_{i^*} \neq \emptyset$, by the f.s. dichotomy (with A^*_i as $\bar{b}$, $A_{<i^*}$ as $\bar{a}$, and c as c), we have that $\mathrm{tp}(A_{i^*}/cA_{<i^*}M)$ is finitely satisfied in M. But this, coupled with $\mathrm{tp}(c/A_{<i^*}M)$ is finitely satisfied in M, would imply $\mathrm{tp}(A_{i^*}c/A_{<i^*}M)$ is finitely satisfied in M, a contradiction.

To finish, we show that for every $j > i^*$, $\mathrm{tp}(A_j/cA_{<j}M)$ is finitely satisfied in M. Again, if this were not the case, the f.s. dichotomy would imply $\mathrm{tp}(cA_j/A_{<j}M)$ is finitely satisfied in M. But then, by Shrinking, we would have $\mathrm{tp}(c/A_{<j}M)$ finitely satisfied in M, contradicting our choice above.

For the 'Moreover' sentence, the only concern is if $\mathrm{tp}(c/A_IM)$ is finitely satisfied in M. But in this case we may take i^* to be the maximal element of I, rather than a new element in $J \backslash I$. □

It is evident that Lemma 3.2 extends to M-f.s. sequences $\langle A_i : i \in I\rangle/C$ over an arbitrary base $C \supseteq M$.

In [?BS, Theorem 4.2.6], the f.s. dichotomy appears as a statement about the behavior of forking rather than non-forking. Namely, forking dependence is totally trivial and transitive on singletons. We may derive similar consequences for dependence from the f.s. dichotomy. This is stated in [?Blum, Corollary 5.22], although missing the necessary condition of full C.

Lemma 3.3. *Suppose T has the f.s. dichotomy, and fix $\bar{a}, \bar{b}, c, M \preceq \mathfrak{C}$.*

1. *If $\mathrm{tp}(\bar{a}/Mc)$ and $\mathrm{tp}(c/M\bar{b})$ are not finitely satisfiable in M, then neither is $\mathrm{tp}(\bar{a}/M\bar{b})$.*
2. *$\mathrm{tp}(\bar{a}/M\bar{b})$ is finitely satisfiable in M if and only if $\mathrm{tp}(a/M\bar{b})$ is finitely satisfied in M for all $a \in \bar{a}$.*
3. *Let $C \supseteq M$ be full. Then $\mathrm{tp}(\bar{a}/C\bar{b})$ is finitely satisfiable in M if and only if $\mathrm{tp}(a/Cb)$ is finitely satisfied in M for all $a \in \bar{a}$, $b \in \bar{b}$.*

Proof. (1) If $\mathrm{tp}(\bar{a}/M\bar{b})$ is finitely satisfiable in M, then by the f.s. dichotomy either $\mathrm{tp}(\bar{a}/M\bar{b}c)$ or $\mathrm{tp}(\bar{a}c/M\bar{b})$ is as well. Shrinking then gives a contradiction.

(2) By induction on $\lg(\bar{a})$. Left to right is immediate, so assume $\mathrm{tp}(a/M\bar{b})$ is finitely satisfied in M for every $a \in \bar{a}$. Write $\bar{a} = \bar{a}'a^*$. By induction we may assume $\mathrm{tp}(\bar{a}'/M\bar{b})$ is finitely satisfied in M. By the f.s. dichotomy, either $\mathrm{tp}(\bar{a}'a^*/M\bar{b})$ is finitely satisfied in M and we are done immediately,

or else $\mathrm{tp}(\bar{a}'/M\bar{b}a^*)$ is finitely satisfied in M, and we finish using transitivity from Fact 2.3.

(3) Left to right is immediate by Shrinking, so assume $\mathrm{tp}(a/Cb)$ is finitely satisfied in M for every $a \in \bar{a}$ and $b \in \bar{b}$. It follows from (2) that $\mathrm{tp}(\bar{a}/Cb)$ is finitely satisfied in M for every $b \in \bar{b}$. To conclude that $\mathrm{tp}(\bar{a}/C\bar{b})$ is finitely satisfied in M, we argue by induction on $\lg(\bar{b})$. Let $\bar{b} = \bar{b}'b^*$, and by induction assume the statement is true for $\bar{b}'$.

By the f.s. dichotomy, either $\mathrm{tp}(\bar{a}/M\bar{b}'b^*)$ or $\mathrm{tp}(\bar{a}'b^*/M\bar{b}')$ is finitely satisfiable in M. In the first case we are finished immediately, and in the second we finish by invoking Lemma 2.13. $\square$

In the stable case, forking dependence is symmetric as well and so yields an equivalence relation on singletons, which is used in [?BS] to decompose models into trees of submodels. In general, the f.s. dichotomy shows finite satisfiability yields a quasi-order on singletons when working over a full $C \supseteq M$. Taking the classes of this quasi-order in order naturally gives an irreducible decomposition of $\mathfrak{C}$ over M in the sense of the next subsection, but we sometimes wish to avoid having to work over a full $C \supseteq M$.

3.1. Decompositions of models.

In this subsection, we characterize the f.s. dichotomy in terms of extending partial decompositions to full decompositions of models.

Definition 3.4 (M-f.s. decomposition)**.** Suppose $X \subseteq \mathfrak{C}$ is any set.

- A *partial M-f.s. decomposition of* X is an M-f.s. sequence $\langle A_i : i \in I\rangle$ with $\bigcup_{i\in I} A_i \subseteq X$.
- An *M-f.s. decomposition of* X is a partial M-f.s. decomposition with $\bigcup_{i\in I} A_i = X$.
- An M-f.s. decomposition of X is *irreducible* if, for every $i \in I$ and for every $a, b \in A_i$, neither $\mathrm{tp}(a/MA_{<i}b)$ nor $\mathrm{tp}(b/MA_{<i}a)$ are finitely satisfied in M.

By iterating Lemma 3.2 for every $c \in X$ for a given set X we obtain:

Lemma 3.5. *Suppose that T has the f.s. dichotomy. For any set $X \subseteq \mathfrak{C}$ and any $C \supseteq M$, any partial M-f.s. decomposition $\langle A_i : i \in I\rangle/C$ of X has a simple extension $\langle A'_j : j \in J\rangle/C$ to an M-f.s. decomposition of X over C. If the sets $\{ A_i : i \in I \}$ were pairwise disjoint, we may choose $\{ A'_j : j \in J \}$ to be pairwise disjoint as well. Furthermore, if $(I, \leq)$ is a well-ordering with a maximum element, we may take $J = I$.*

Proposition 3.6. *T has the f.s. dichotomy if and only if for all models $M, N \preceq \mathfrak{C}$ (we do not require $M \subseteq N$) every partial M-f.s. decomposition $\langle A_i : i \in I\rangle$ of N has an irreducible M-f.s. decomposition of N extending it.*

Proof. ($\Leftarrow$) Suppose partial decompositions extend, and let $\bar{a}, \bar{b}, c \in \mathfrak{C}, M \prec \mathfrak{C}$ with $\mathrm{tp}(\bar{b}/\bar{a})$ finitely satisfiable in M, and let $N \prec \mathfrak{C}$ with $\bar{a}, \bar{b}, c \in N$. So

$\langle \bar{a}, \bar{b} \rangle$ is an M-f.s. sequence, and can be extended to an M-f.s. decomposition of N. After Shrinking, we obtain the conclusion of the f.s.-dichotomy.

($\Rightarrow$) Given a partial M-f.s. decomposition $\langle A_i : i \in I \rangle$ of N, apply Lemma 3.5 to get a simple extension $\langle B_k : k \in K \rangle$ that is an M-f.s. decomposition of N. By Zorn's Lemma, it will suffice to show that if there is $k \in K$ with $a, b \in B_k$ and $\mathrm{tp}(b/B_{<k}a)$ is finitely satisfiable in M, then we may blow up the sequence so that a and b are in distinct parts. Given such $a, b \in B_k$, we have $\langle a, b \rangle / B_{<k}$ is a partial M-f.s. decomposition of B_k over $B_{<k}$. Extending this to a full decomposition of B_k and then applying Lemma 2.8 to prepend $B_{<k}$ and append $B_{>k}$ gives the result. $\square$

Remark 3.7. We could do the above proof over some full $C \supseteq M$ to obtain an irreducible M-f.s. decomposition that is also an order-congruence.

We note that at the end of [?BS], Baldwin and Shelah conjecture that models of monadically NIP theories should admit tree decompositions like those they describe for monadically stable theories, but with order-congruences in place of full congruences.

3.2. **Preserving indiscernibility.** We begin with some definitions. The definition of dp-minimality given here may be non-standard, but it is proven equivalent to the usual definition with Fact 2.10 of [?DGL].

Definition 3.8 (Indiscernible-triviality and dp-minimality)**.** The first definition is meant to recall trivial forking.

- T has *indiscernible-triviality* if for every infinite indiscernible sequence $\mathcal{I}$ and every set B of parameters, if $\mathcal{I}$ is indiscernible over each $b \in B$ then $\mathcal{I}$ is indiscernible over B.
- T is *dp-minimal* if, for all indiscernible sequences $\mathcal{I} = \langle \bar{a}_i : i \in I \rangle$ over any set C, every $b \in \mathfrak{C}$ induces a finite partition of the index set into convex pieces $\mathcal{I} = I_1 \ll I_2 \ll \cdots \ll I_n$, with at most two I_j infinite and every $\mathcal{I}_j = \langle \bar{a}_i : i \in I_j \rangle$ is indiscernible over Cb.

As mentioned in the introduction, the notion of a theory admitting coding was the central dividing line of [?BS]. We weaken the definition here to allow the sequences to consist of tuples. Note that even the theory of equality would admit the further weakening of also allowing C to consist of tuples.

Definition 3.9 (Admits coding (on tuples))**.** A theory T *admits coding on tuples* if there is a formula $\phi(\bar{x}, \bar{y}, z)$ (with parameters $\bar{d}$), sequences $\mathcal{I} = \langle \bar{a}_i : i \in I \rangle, \mathcal{J} = \langle \bar{b}_j : j \in J \rangle$, indexed by countable, dense orderings $(I, \leq), (J, \leq)$, respectively, and a set $\{ c_{i,j} \mid i \in I, j \in J \}$, such that $\mathcal{I}$ is indiscernible over $\bigcup \mathcal{J} \cup \bar{d}$, $\mathcal{J}$ is indiscernible over $\bigcup \mathcal{I} \cup \bar{d}$, and $\mathfrak{C} \models \phi(\bar{a}_i, \bar{b}_j, c_{k,l}) \iff (i, j) = (k, l)$.

We call $\langle \bar{a}_i : i \in I \rangle, \langle \bar{b}_j : j \in J \rangle, \{ c_{i,j} \mid i \in I, j \in J \}, \phi(\bar{x}, \bar{y}, z)$ a *tuple-coding configuration*, and let $A = \bigcup \{ \bar{a}_i : i \in I \}, B = \bigcup \{ \bar{b}_j : j \in J \}$ and $C = \{ c_{i,j} \mid i \in I, j \in J \}$.

T *admits coding* if we may take $\mathcal{I}$ and $\mathcal{J}$ to be sequences of singletons.

A convenient variant for this subsection is a *joined tuple-coding configuration*, which consists of a formula (with parameters) $\phi(\bar{x}, \bar{y}, z)$, a sequence $\langle \bar{a}_i : i \in I \rangle$ indiscernible over the parameters of ϕ, indexed by an infinite linear order $(I, \leq)$, and a set $\{ c_{i,j} \mid i < j \in I \}$ such that for $i < j$, $\mathfrak{C} \models \phi(\bar{a}_i, \bar{a}_j, c_{k,l}) \iff (i, j) = (k, l)$. Given a joined tuple-coding configuration, indexed by a countable, dense $(I, \leq)$, we may construct a tuple-coding configuration by keeping $\phi(\bar{x}, \bar{y}, z)$ fixed, choosing open intervals $I', J' \subseteq I$ with $I' \ll J'$, and letting $\mathcal{I}' = \langle \bar{a}_i : i \in I' \rangle$ and $\mathcal{J}' = \langle \bar{a}_j : j \in J' \rangle$. Conversely, given a tuple-coding configuration with $I = J$, we may construct a joined tuple-coding configuration by considering the indiscernible sequence whose elements are $\langle \bar{a}_i \bar{b}_i : i \in I \rangle$, restricting C to elements $c_{i,j}$ with $i < j$, and replacing ϕ by $\phi^*(\bar{x}\bar{x}', \bar{y}\bar{y}', z) := \phi(\bar{x}, \bar{y}', z)$.

The following configuration appears in [?Hanf, Part II Lemma 2.2], and will appear as an intermediate between a failure of the f.s. dichotomy and a tuple-coding configuration.

Definition 3.10. A *pre-coding configuration* consists of a $\phi(\bar{x}, \bar{y}, z)$ with parameters and a sequence $\mathcal{I} = \langle \bar{d}_i : i \in \mathbb{Q} \rangle$, indiscernible over the parameters of ϕ, such that for some (equivalently, for every) $s < t$ from $\mathbb{Q}$, there is $c \in \mathfrak{C}$ such that

(1) $\mathfrak{C} \models \phi(\bar{d}_s, \bar{d}_t, c)$;
(2) $\mathfrak{C} \models \neg\phi(\bar{d}_s, \bar{d}_v, c)$ for all $v > t$; and
(3) $\mathfrak{C} \models \neg\phi(\bar{d}_u, \bar{d}_t, c)$ for all $u < s$.

We show the equivalence of the existence of these notions with the proposition below. The proof of (4) $\Rightarrow$ (1) in the following is essentially from [?Hanf, Part II Lemma 2.3], while (3) $\Rightarrow$ (4) is based on [?ST, Lemma C.1]. The idea of (4) $\Rightarrow$ (1) is that when working over a full $D \supseteq M$, types have a unique "generic" extension by Lemma 2.12. In a failure of the f.s. dichotomy, the extension of $\mathrm{tp}(c/D)$ to $\mathrm{tp}(c/Dab)$ is non-generic, and so c can in some sense pick out a and b from a suitable sequence.

Proposition 3.11. *The following are equivalent for any theory T.*

(1) T has the f.s. dichotomy.
(2) T is dp-minimal and has indiscernible-triviality.
(3) T does not admit coding on tuples.
(4) T does not have a pre-coding configuration.

Proof. (1) $\Rightarrow$ (2): Suppose T has the f.s. dichotomy. We begin with showing T is dp-minimal. Choose an indiscernible $\mathcal{I} = \langle \bar{a}_i : i \in I \rangle$ over a set D and any element $b \in \mathfrak{C}$. Applying Lemma 2.20 to $\mathcal{I}$ (in the theory T_D naming constants for each $d \in D$) choose a model $M \supseteq D$ and a full set $C \supseteq M$ such that $\mathcal{I}$ is both indiscernible over C and an M-f.s. sequence over C. As in the proof of Lemma 3.2, choose a maximal initial segment $I_0 \subseteq I$ such that $\mathrm{tp}(b/A_{I_0}C)$ is finitely satisfied in M. If $(I \setminus I_0)$ has a minimal element i^*, let $I_1 = (I \setminus (I_0 \cup \{ i^* \}))$, and let $I_1 = I \setminus I_0$ otherwise. As C is full, in either case we have an order-congruence of $\mathcal{I} \cup \{ b \}$, so both I_0 and I_1 are indiscernible over Cb, which suffices.

Next, we show indiscernible-triviality. We may assume $\mathcal{I} = \langle \bar{a}_i : i \in \mathbb{Q} \rangle$ is ordered by $(\mathbb{Q}, \leq)$. The argument here is more involved, as given an infinite, indiscernible $\mathcal{I}$ and a set B over which $\mathcal{I}$ is indiscernible over each $b \in B$, we cannot apply Lemma 2.20 and maintain the indiscernibility over each $b \in B$. However, the proof of Lemma 2.18 allows us to find a model M such that $\mathcal{I}$ is an M-f.s. sequence and is indiscernible over Mb for every $b \in B$. Now, call an element $b \in B$ *high* if $\mathrm{tp}(b/MI)$ is finitely satisfied in M. By indiscernibility, if $\mathrm{tp}(b/MA_{<i})$ is finitely satisfied in M for any i, then b is high. Let $B_1 \subseteq B$ denote the set of high elements, and let $B_0 := B \setminus B_1$ be the *low* (i.e., not high) elements of B. As T satisfies the f.s. dichotomy, Lemma 3.5 implies $\mathcal{I} = \langle \bar{a}_i : i \in \mathbb{Q} \rangle$ has a simple extension (Definition 2.7) to an M-f.s. sequence with universe $(\bigcup \mathcal{I})B$. Moreover, any such simple extension will condense to $\langle B_0, \mathcal{I}, B_1 \rangle$.

Using this, we argue that $\mathcal{I}$ is indiscernible over MB in two steps. First, we argue that $\mathcal{I}$ is indiscernible over MB_0. To see this, fix $i < j$ from $\mathbb{Q}$ and let $p_i = \mathrm{tp}(\bar{a}_i/A_{<i}MB_0)$. From the previous paragraph, $\mathcal{I}$ is an M-f.s. sequence over MB_0. So p_i does not split over M, and so by Lemma 2.21 it suffices to prove that $\bar{a}_j$ realizes p_i. Choose any $\phi(\bar{x}, \bar{b}, \bar{m}) \in p_i$ with $\bar{m}$ from M and $\bar{b}$ from B_0. To see that $\mathfrak{C} \models \phi(\bar{a}_j, \bar{b}, \bar{m})$, choose an automorphism $\sigma \in Aut(\mathfrak{C})$ fixing M and an initial segment $\langle \bar{a}_i : i \in I_0 \rangle$ pointwise that induces an order-preserving permutation of $\mathcal{I}$ with $\sigma(\bar{a}_i) = \bar{a}_j$. Clearly, $\mathfrak{C} \models \phi(\bar{a}_j, \sigma(\bar{b}), \bar{m})$. It is easily seen that for every singleton $b' \in \sigma(\bar{b})$, $\mathcal{I}$ is indiscernible over Mb' and, as σ fixes A_{I_0} pointwise, b' is also low. Thus, any simple extension to $(\bigcup \mathcal{I})MB_0\sigma(\bar{b})$ will condense to $\langle B_0\sigma(\bar{b})\rangle \frown \langle \bar{a}_i : i \in \mathbb{Q} \rangle \frown \langle B_1 \rangle$. In particular $\mathrm{tp}(\bar{a}_j/M\bar{b}\sigma(\bar{b}))$ is finitely satisfied in M. Thus, if $\mathfrak{C} \models \neg\phi(\bar{a}_j, \bar{b}, \bar{m})$, by finite satisfiability there would be $\bar{n}$ from M such that $\mathfrak{C} \models \neg\phi(\bar{n}, \bar{b}, \bar{m}) \wedge \phi(\bar{n}, \sigma(\bar{b}), \bar{m})$, which is impossible since σ fixes M pointwise. Thus, $\mathcal{I}$ is indiscernible over MB_0.

Finally, to see that $\mathcal{I}$ is indiscernible over MB_0B_1, choose any $i_1 < \ldots i_k$, $j_1 < \cdots < j_k$ from $\mathbb{Q}$, $\bar{b}$ from MB_0, and $\bar{c}$ from B_1 and assume by way of contradiction that $\mathfrak{C} \models \psi(\bar{a}_{i_1}, \ldots, \bar{a}_{i_k}, \bar{b}, \bar{c}) \wedge \neg\psi(\bar{a}_{j_1}, \ldots, \bar{a}_{j_k}, \bar{b}, \bar{c})$. Recall $\langle B_0, \mathcal{I}, B_1 \rangle$ is an M-f.s. sequence, so $\mathrm{tp}(\bar{c}/M(\bigcup \mathcal{I})B_0)$ is finitely satisfied in M, and so the same formula is true with some $\bar{m}$ from M replacing $\bar{c}$. But this contradicts that $\mathcal{I}$ is indiscernible over MB_0. Thus, $\mathcal{I}$ is indiscernible over MB.

(2) $\Rightarrow$ (3): Assume (2) holds, but (3) fails, so there is a joined tuple-coding configuration $\langle \bar{a}_i : i \in I \rangle, \{ c_{i,j} \mid i < j \in I \}, \phi(\bar{x}, \bar{y}, z)$ with $(I, \leq)$ countable, dense. By naming constants, we may assume ϕ has no parameters. Choose $i < j$ from I. By dp-minimality, $(I, \leq)$ is partitioned into finitely many convex pieces, indiscernible over $c_{i,j}$, with at most two pieces infinite.

If $\bar{a}_i, \bar{a}_j$ are in the same convex piece, then taking $i < k < j$ we get $\phi(\bar{a}_k, \bar{a}_j, c_{i,j})$, contradicting our configuration. So suppose $\bar{a}_i$ and $\bar{a}_j$ are in different pieces. Then one of the pieces must be infinite, so by symmetry suppose the piece I' containing $\bar{a}_i$ is. By indiscernible-triviality $(\bar{a}_i : i \in I')$ is indiscernible over $\bar{a}_j c_{i,j}$. But then picking some $k \in I' \setminus \{ i \}$ again gives $\phi(\bar{a}_k, \bar{a}_j, c_{i,j})$.

(3) $\Rightarrow$ (4): Assume (4) fails, as witnessed by $\phi(\bar{x}, \bar{y}, z)$, $\langle \bar{a}_i : i \in \mathbb{Q} \rangle$, and $\{ c_{i,j} : i < j \in \mathbb{Q} \}$. Define a new sequence $(\bar{b}_i : i \in \mathbb{Z})$ where $\bar{b}_i = \bar{a}_{3i-1}\bar{a}_{3i}\bar{a}_{3i+1}$. Let

$$\psi(\bar{x}_-\bar{x}\bar{x}_+, \bar{y}_-\bar{y}\bar{y}_+, z) := \phi(\bar{x}, \bar{y}, z) \wedge \neg\phi(\bar{x}_-, \bar{y}, z) \wedge \neg\phi(\bar{x}, \bar{y}_+, z)$$

Also let $d_{i,j} = c_{3i,3j}$. It is easily verified that $(\bar{b}_i : i \in \mathbb{Z}), \{ d_{i,j} \mid i < j \in \mathbb{Z} \}, \psi$ is a joined tuple-coding configuration. That T admits coding on tuples now follows by the remarks following Definition 3.9.

(4) $\Rightarrow$ (1): Assume that $\bar{a}$, $\bar{b}$, M, and c form a counterexample to the f.s. dichotomy, i.e., $\mathrm{tp}(\bar{b}/M\bar{a})$ is finitely satisfied in M, but neither $\mathrm{tp}(\bar{b}c/M\bar{a})$ nor $\mathrm{tp}(\bar{b}/M\bar{a}c)$ are. As $\langle \bar{a}, \bar{b} \rangle$ is an M-f.s. sequence, by Proposition 2.10 choose a full $D \supseteq M$ such that $\langle \bar{a}, \bar{b} \rangle / D$ is an M-f.s. sequence over D. Note that by transitivity, we also have $\mathrm{tp}(\bar{a}\bar{b}/D)$ finitely satisfied in M.

Claim. There is c' such that $\bar{a}\bar{b}c' \equiv_M \bar{a}\bar{b}c$ with $\mathrm{tp}(\bar{a}\bar{b}c'/D)$ finitely satisfied in M.

Proof of Claim. We first argue that every finite conjunction of formulas from $\mathrm{tp}(\bar{a}\bar{b}/D) \cup \mathrm{tp}(\bar{a}\bar{b}c/D)$ is satisfied in M. To see this choose $\phi(\bar{x}, \bar{y}, \bar{d}) \in \mathrm{tp}(\bar{a}\bar{b}/D)$ and $\psi(\bar{x}, \bar{y}, z) \in \mathrm{tp}(\bar{a}\bar{b}c/M)$ (ϕ and ψ may also have hidden parameters from M) and we will show that $\phi(\bar{x}, \bar{y}, \bar{d}) \wedge \psi(\bar{x}, \bar{y}, z)$ has a solution in M. Let $\theta(\bar{x}, \bar{y}, \bar{d}) := \phi(\bar{x}, \bar{y}, \bar{d}) \wedge \exists z \psi(\bar{x}, \bar{y}, z)$. As $\mathrm{tp}(\bar{a}\bar{b}/D)$ is finitely satisfied in M, $\theta(\bar{m}, \bar{n}, \bar{d})$ holds for some $\bar{m}, \bar{n}$ from M. Thus, as $M \preceq \mathfrak{C}$ and $\mathfrak{C} \models \exists z \psi(\bar{m}, \bar{n}, z)$, there is $k \in M$ such that $\psi(\bar{m}, \bar{n}, k)$ holds, which suffices.

Thus, by Fact 2.3(2), there is a complete type $p(\bar{x}, \bar{y}, z) \in S(D)$ extending $\mathrm{tp}(\bar{a}\bar{b}/D) \cup \mathrm{tp}(\bar{a}\bar{b}c/D)$ that is finitely satisfied in M. As $\mathrm{tp}(\bar{a}\bar{b}/D) \subseteq p$, there is an element c' so that $(\bar{a}, \bar{b}, c')$ realizes p, proving the claim. $\Diamond$

By Lemma 2.23, choose an M-f.s. sequence $\langle \bar{a}_i\bar{b}_i c_i : i \in \mathbb{Q} \rangle$ over D of realizations of p that is indiscernible over D. Fix $s < t$ from $\mathbb{Q}$. Note that since $\bar{a}, \bar{b}, c, M$ witness the failure of the f.s. dichotomy, neither $\mathrm{tp}(\bar{b}_s c_s/D\bar{a}_s)$ nor $\mathrm{tp}(\bar{b}_s/D\bar{a}_s c_s)$ are finitely satisfied in M. As notation, let $\bar{a}_{<s} = \bigcup \{ \bar{a}_i : i < s \}$ and $\bar{b}_{>t} = \bigcup \{ \bar{b}_j : j > t \}$. Now, by Shrinking and Condensation,

$$\langle \bar{a}_{<s}, (\bar{a}_s\bar{b}_s), \bar{b}_t, \bar{b}_{>t} \rangle \text{ is an } M\text{-f.s. sequence of length 4 over } D.$$

As $\mathrm{tp}(\bar{b}_s/D\bar{a}_s)$ is finitely satisfied in D and D is full, by Lemma 2.13 $\langle \bar{a}_{<s}, \bar{a}_s, \bar{b}_s \rangle$ is an M-f.s. sequence over D, and so by Lemma 2.8,

$$\langle \bar{a}_{<s}, \bar{a}_s, \bar{b}_s, \bar{b}_t, \bar{b}_{>t} \rangle \text{ is an } M\text{-f.s. sequence of length 5 over } D.$$

As this sequence is an order-congruence, it follows that $\mathrm{tp}(\bar{b}_s/D\bar{a}_{<s}\bar{a}_s\bar{b}_{>t}) = \mathrm{tp}(\bar{b}_t/D\bar{a}_{<s}\bar{a}_s\bar{b}_{>t})$, so we can choose c^* such that

$$\bar{b}_s c_s \equiv_{D\bar{a}_{<s}\bar{a}_s\bar{b}_{>t}} \bar{b}_t c^*$$

Thus, since $\mathrm{tp}(\bar{b}_s/D\bar{a}_s c_s)$ is not finitely satisfied in M, neither is $\mathrm{tp}(\bar{b}_t/D\bar{a}_s c^*)$. By contrast, for $j, j' > t$, as both $\mathrm{tp}(\bar{b}_j/D\bar{a}_s c_s)$ and $\mathrm{tp}(\bar{b}_{j'}/D\bar{a}_s c_s)$ are finitely satisfied in M and are equal by Proposition 2.17, the same is true of $\mathrm{tp}(\bar{b}_j/D\bar{a}_s c^*) = \mathrm{tp}(\bar{b}_{j'}/D\bar{a}_s c^*)$. Dually, since $\mathrm{tp}(\bar{b}_s c_s/D\bar{a}_s)$ is not finitely satisfied in M, neither is $\mathrm{tp}(\bar{b}_s c^*/D\bar{a}_s)$; and because $\mathrm{tp}(\bar{b}_s c_s/D\bar{a}_i)$ is finitely satisfied in M for every $i < s$, we also conclude $\mathrm{tp}(\bar{b}_s c^* \bar{a}_i/D) = \mathrm{tp}(\bar{b}_s c^* \bar{a}_{i'}/D)$ for all $i, i' < s$ by Proposition 2.17.

Finally, choose $\rho_1(\bar{x}, \bar{y}, z), \rho_2(\bar{x}, \bar{y}, z) \in \mathrm{tp}(\bar{a}_s, \bar{b}_t, c^*/D)$ such that neither $\rho_1(\bar{a}_s, \bar{y}, c^*)$ nor $\rho_2(\bar{a}_s, \bar{y}, z)$ has realizations in M. Then, letting $\bar{d}_i := \bar{a}_i\bar{b}_i$ for each $i \in \mathbb{Q}$, $\mathcal{I} = \langle \bar{d}_i : i \in \mathbb{Q}\rangle$ is a pre-coding configuration with respect to $\rho_1 \wedge \rho_2$ and c^*. ☐

Remark 3.12. The following observation will be useful in Section 5. A *tidy* pre-coding configuration $\mathcal{I} = \langle \bar{d}_i : i \in \mathbb{Q}\rangle, \{ c_{s,t} : s < t \in \mathbb{Q} \}, \phi(\bar{x}, \bar{y}, z)$ is one where $\mathfrak{C} \models \neg\phi(\bar{d}_i, \bar{d}_j, e)$ for every $i < j$ and $e \in \bigcup \mathcal{I}$. The pre-coding configuration constructed in (4) $\Rightarrow$ (1) is tidy, since, choosing M so $\mathcal{I}$ is an M-f.s. sequence, $\mathfrak{C} \models \phi(\bar{d}_i, \bar{d}_j, e)$ implies $\mathrm{tp}(\bar{d}_j/\bar{d}_i e)$ and $\mathrm{tp}(\bar{d}_j e/\bar{d}_i)$ are not finitely satisfiable in M. But if $e \in \bar{d}_k$ for $k \leq i$ then the former type is finitely satisfiable in M, and if $e \in \bar{d}_k$ for $k > i$, then the latter type is.

The tidiness property extends to the joined tuple-coding configuration constructed in (3) $\Rightarrow$ (4) and so ultimately to the tuple-coding configuration as well. That is, from a failure of the f.s. dichotomy, we construct a tuple-coding configuration $\mathcal{I}, \mathcal{J}, C, \phi$ with $\mathfrak{C} \models \neg\phi(\bar{a}_i, \bar{b}_j, e)$ for every $\bar{a}_i \in \mathcal{I}, \bar{b}_j \in \mathcal{J}$, and $e \in \bigcup \mathcal{I} \cup \bigcup \mathcal{J}$.

4. The main theorem

We recall the main theorem from the introduction. Note that whereas Clauses (1) and (2) discuss monadic expansions of T, Clauses (3)-(6) are all statements about T itself.

Theorem 4.1. *The following are equivalent for a complete theory T with an infinite model.*

(1) T is monadically NIP.
(2) No monadic expansion of T admits coding.
(3) T does not admit coding on tuples.
(4) T has the f.s. dichotomy.
(5) For all $M^* \models T$ and $M, N \preceq M^*$, every partial M-f.s. decomposition of N extends to an (irreducible) M-f.s. decomposition of N.
(6) T is dp-minimal and has indiscernible-triviality.

The equivalences of (3)-(6) are by Proposition 3.6 and Proposition 3.11. We note that (1) $\Rightarrow$ (2) is easy: Choose a monadic expansion $\mathfrak{C}^*$ that admits coding, say via an L-formula $\phi(x, y, z)$ defining a bijection from the countable sets $A \times B \to C$. By adding a new unary predicate for a suitable $C_0 \subseteq C$, the formula $\psi(x, y) := \exists z \in C_0 \phi(x, y, z)$ can define the edge relation of an arbitrary bipartite graph on $A \times B$, and in particular of the generic bipartite graph. Thus, T is not monadically NIP.

Thus, it remains to prove that (4) $\Rightarrow$ (1) and (2) $\Rightarrow$ (3), which are proved in the next two subsections.

4.1. **If T has the f.s. dichotomy, then T is monadically NIP.** The type-counting argument in this section is somewhat similar to that in [?Blum], showing that monadic NIP corresponds to the dichotomy of unbounded partition width versus partition width at most $\beth_2(\aleph_0)$. Both arguments use the tools from Sections 2 and 3 to decompose the model and count the types realized in a part of the decomposition over its complement. However, while Blumensath decomposes the model into a large binary tree, our decomposition takes a single step.

Definition 4.2. Suppose $\mathcal{I} = \{ \bar{a}_i : i \in I \}$ is any sequence of pairwise disjoint tuples and suppose $N \supseteq \bigcup \mathcal{I}$ is any model. An *$\mathcal{I}$-partition of N* is any partition $N = \bigsqcup \{ A_i : i \in I \}$ with $\bar{a}_i \subseteq A_i$ for each $i \in I$.

Definition 4.3. For any $N \preceq \mathfrak{C}$ and $A \subseteq N$, let $\mathrm{rtp}(N, A)$ denote the number of complete types over A *realized* by tuples in $(N \setminus A)^{<\omega}$.

We will be primarily interested in the case where A is very large, and $\mathrm{rtp}(N, A)$ is significantly smaller than $|A|$. The following lemma is similar to Lemma 2.15, removing the requirement that the partition is convex but adding a finiteness condition.

For the rest of this section, recall the notation $A_J = \bigcup_{j \in J} A_j$ from the first part of Definition 2.4.

Lemma 4.4. *If T has the f.s. dichotomy, then for every well-ordering $(I, \leq)$ with a maximum element, for every indiscernible sequence $\mathcal{I} = \{ \bar{a}_i : i \in I \}$ of pairwise disjoint tuples and every $N \supseteq \bigcup \mathcal{I}$, there is an $\mathcal{I}$-partition $\{ A_i : i \in I \}$ of N such that $\mathrm{rtp}(N, A_J) \leq \beth_2(|T|)$ for every finite $J \subseteq I$.*

Proof. By Lemma 2.20, choose a model M of size $|T|$ and a full $C \supseteq M$ with $|C| \leq 2^{|T|}$ for which $\langle \bar{a}_i : i \in I \rangle / C$ is an M-f.s. sequence over C. (Note that N might not contain M.) By Lemma 3.5 choose a simple extension $\langle A_i : i \in I \rangle$ of $\langle \bar{a}_i : i \in I \rangle$ with $\bigsqcup \{ A_i : i \in I \} = N$. Thus, $\{ A_i : i \in I \}$ is an $\mathcal{I}$-partition of N. For a given finite $J \subseteq I$ and $\bar{n} \in N \setminus A_J$, let $\bar{n} \subset A_{i_1} \cup \cdots \cup A_{i_k}$ and let $\bar{n}_{i_j} = \bar{n} \cap A_{i_j}$. As $\langle A_i : i \in I \rangle / C$ is an order-congruence over C by Lemma 2.17, $\mathrm{tp}(\bar{n}/CA_J)$ is determined by $\{ \mathrm{tp}(\bar{n}_{i_j}/C) : 1 \leq j \leq k \}$ and the order type of the finite set $\{ i_1, \ldots, i_k \} \cup J$. There are only finitely many such order types, and as $|C| \leq 2^{|T|}$, there are at most $\beth_2(|T|)$ complete types over C. So $\mathrm{rtp}(N, A_J) \leq \beth_2(|T|)$ for every finite $J \subseteq I$. $\square$

On the other hand, if a theory T has the independence property, then no uniform bound can exist.

Lemma 4.5. *Suppose that T has IP, as witnessed by $\phi(\bar{x}, y)$ with $\lg(\bar{x}) = n$ and $\lg(y) = 1$. For every $\lambda > 2^{|T|}$ there is an order-indiscernible $\mathcal{I} = (a_i :$*

$i \leq \lambda$), and a model $N \supseteq \mathcal{I}$ such that for every $\mathcal{I}$-partition $\langle A_i : i \leq \lambda \rangle$ of N, $\mathrm{rtp}(N, A_{I_0}) \geq \lambda$ *for some finite $I_0 \subseteq (\lambda + 1)$.*

Proof. In the monster model, choose an order-indiscernible $\mathcal{I} = (a_i : i \leq \lambda)$ that is shattered, i.e., there is a set $Y = \{ \bar{b}_s : s \in \mathcal{P}(\lambda) \}$ such that $\phi(\bar{b}_s, a_i)$ holds if and only if $i \in s$. Note that for distinct $\bar{b}, \bar{b}' \in Y$, there is some $a_i \in I$ such that $\mathrm{tp}_\phi(\bar{b}/a_i) \neq \mathrm{tp}_\phi(\bar{b}'/a_i)$. Let N be any model containing $\mathcal{I} \cup \bigcup Y$ and let $\langle A_i : i \leq \lambda \rangle$ be any $\mathcal{I}$-partition of N. As $|\mathcal{I}| = \lambda$, while $|Y| = 2^\lambda$, by applying the pigeon-hole principle n times (one for each coordinate of $\bar{b}$) one obtains $Y' \subseteq Y$, also of size 2^λ, and a finite $I_0 \subseteq I$ such that $\bar{b} \in (A_{I_0})^n$ for each $\bar{b} \in Y'$. As $\lambda > 2^{|T|}$ and there are at most $2^{|T|}$ types over I_0, we can find $Y^* \subseteq Y'$ of size 2^λ such that $\mathrm{tp}(\bar{b}/I_0)$ is constant among $\bar{b} \in Y^*$. It follows that $\mathrm{tp}(\bar{b}/(I - I_0)) \neq \mathrm{tp}(\bar{b}'/(I - I_0))$ for distinct $\bar{b}, \bar{b}' \in Y^*$. As $Y^* \subseteq (A_{I_0})^n$, it follows that $\mathrm{rtp}(N, A_{I_0}) \geq \lambda$. □

To show that the behaviors of Lemma 4.4 and Lemma 4.5 cannot co-exist, we get an upper bound on the number of types realized in a finite monadic expansion. Such a bound is easy for quantifier-free types, and the next lemma inductively steps it up to a bound on all types. The following two lemmas make no assumptions about T.

For each $k \in \omega$, define an equivalence relation $\sim_k$ on $(N \setminus A)^{<\omega}$ by: $\bar{a} \sim_k \bar{b}$ if and only if $\mathrm{lg}(\bar{a}) = \mathrm{lg}(\bar{b})$ and $\mathrm{tp}_\phi(\bar{a}/A) = \mathrm{tp}_\phi(\bar{b}/A)$ for every formula $\phi(\bar{z})$ of quantifier depth at most k. Clearly, $\mathrm{tp}(\bar{a}/A) = \mathrm{tp}(\bar{b})$ if and only if $\bar{a} \sim_k \bar{b}$ for every k. To get an upper bound on $\mathrm{rtp}(N, A)$, for each $k \in \omega$, let $r_k(N, A) = |(N \setminus A)^{<\omega} / \sim_k|$.

Lemma 4.6. *For any $N \preceq \mathfrak{C}$, $A \subseteq N$, and $k \in \omega$, $r_{k+1}(N, A) \leq 2^{r_k(N,A)}$. Thus,* $\mathrm{rtp}(N, A) \leq \beth_{\omega+1}(r_0(N, A))$.

Proof. The second sentence follows from the first as $\mathrm{tp}(\bar{a}/A) = \mathrm{tp}(\bar{b}/A)$ if and only if $\bar{a} \sim_k \bar{b}$ for every k. For the first sentence, we give an alternate formulation of $\sim_k$ to make counting easier. For each $k \in \omega$, let E_k be the equivalence relation on $(N \setminus A)^{<\omega}$ given by:

- $E_0(\bar{a}, \bar{b})$ if and only if $\mathrm{lg}(\bar{a}) = \mathrm{lg}(\bar{b})$ and $\mathrm{qftp}(\bar{a}/A) = \mathrm{qftp}(\bar{b}/A)$; and
- $E_{k+1}(\bar{a}, \bar{b})$ if and only if $E_k(\bar{a}, \bar{b})$ and, for every $c \in (N \setminus A)$, there is $d \in (N \setminus A)$ such that $E_k(\bar{a}c, \bar{b}d)$, and vice-versa,

For each k, let $c(k) := |(N \setminus A)^{<\omega} / E_k|$. It is clear that $c(0) = r_0(N, A)$ and by the definition of E_{k+1} we have $c(k+1) \leq 2^{c(k)}$ for each k, so the lemma follows from the fact that $E_k(\bar{a}, \bar{b})$ if and only if $\bar{a} \sim_k \bar{b}$, whose verification amounts to proving the following claim.

Claim. If the quantifier depth of $\phi(\bar{z})$ is at most k, then for all partitions $\bar{z} = \bar{x}\bar{y}$, for all $\bar{e} \in A^{\mathrm{lg}(\bar{y})}$, and for all $\bar{a}, \bar{b} \in (N \setminus A)^{\mathrm{lg}(\bar{x})}$, if $E_k(\bar{a}, \bar{b})$, then $N \models \phi(\bar{a}, \bar{e}) \leftrightarrow \phi(\bar{b}, \bar{e})$.

Proof of Claim. By induction on k. Say $\psi(\bar{z}) := \exists w \phi(w, \bar{z})$ is chosen with the quantifier depth of ϕ is at most k. Fix a partition $\bar{z} = \bar{x}\bar{y}$ and choose

$\bar{e} \in A^{\lg(\bar{y})}$, $\bar{a}, \bar{b} \in (N \setminus A)^{\lg(\bar{x})}$ with $E_{k+1}(\bar{a}, \bar{b})$. Assume $N \models \exists w \phi(w, \bar{a}, \bar{e})$. There are two cases. If there is some $h \in A$ such that $N \models \phi(h, \bar{a}, \bar{e})$, then $N \models \phi(h, \bar{b}, \bar{e})$ by the inductive hypothesis. On the other hand, if there is $c \in (N \setminus A)$ such that $N \models \phi(c, \bar{a}, \bar{e})$, use $E_{k+1}(\bar{a}, \bar{b})$ to find $d \in (N \setminus A)$ such that $E_k(\bar{a}c, \bar{b}d)$. Thus, the inductive hypothesis implies $N \models \phi(d, \bar{b}, \bar{e})$, so again $N \models \psi(\bar{b}, \bar{e})$. $\Diamond$

$\square$

The following transfer result is the point of the previous lemma. Again, it will be used when $\mathrm{rtp}(N^+, A)$ is significantly smaller than $|A|$.

Lemma 4.7. *Let $N \supseteq A$ be any model and let $N^+ = (N, U_1, \ldots, U_k)$ be any expansion of N by finitely many unary predicates. Then* $\mathrm{rtp}(N^+, A) \leq \beth_{\omega+1}(\mathrm{rtp}(N, A))$.

Proof. For each n, expanding by k unary predicates can increase the number of quantifier-free n-types by at most a finite factor, i.e. 2^k, so $r_0(N^+, A) = r_0(N, A) \leq \mathrm{rtp}(N, A)$. The result now follows from Lemma 4.6. $\square$

Finally, we combine the lemmas above to obtain the goal of this subsection.

Proposition 4.8. *If T has the f.s. dichotomy, then T is monadically NIP.*

Proof. By way of contradiction assume that T is not monadically NIP, but has the f.s. dichotomy. Let T^+ be an expansion by finitely many unary predicates that has IP. Choose a cardinal $\lambda > \beth_{\omega+1}(|T|)$. Let $N^+ \models T^+$ with $N^+ \supseteq \mathcal{I} = (a_i : i \leq \lambda)$ as in Lemma 4.5, so for any $\mathcal{I}$-partition of N^+ there is $I_0 \subseteq (\lambda + 1)$ with $\mathrm{rtp}(N^+, A_{I_0}) > \beth_{\omega+1}(|T|)$.

Let N be the L-reduct of N^+. As $\mathcal{I}$ remains L-order-indiscernible, and T has the f.s. dichotomy, choose an $\mathcal{I}$-partition $\langle A_i : I \leq \lambda \rangle$ of N as in Lemma 4.4, so $\mathrm{rtp}(N, A_J) \leq \beth_2(|T|)$ for every $J \subseteq (\lambda + 1)$. Since N^+ is a unary expansion of N, $\mathrm{rtp}(N^+, A_J) \leq \beth_{\omega+1}(|T|)$ for every $J \subseteq (\lambda + 1)$, by Lemma 4.7. This this contradicts our ability to find an $I_0 \subseteq (\lambda + 1)$ from the previous paragraph for the chosen $\mathcal{I}$-partition of N^+. $\square$

Lemma 2.15 and the arguments in this subsection seem to indicate that, for a generalization of the structural graph-theoretic notion of neighborhood-width [?Gur] similar to Blumensath's generalization of clique-width [?Blum], monadic NIP should correspond to a dichotomy between bounded and unbounded neighborhood-width.

4.2. **From coding on tuples to coding on singletons.** This subsection provides the final step, (2) $\Rightarrow$ (3), in proving Theorem 4.1 by showing that if T admits coding on tuples, then some monadic expansion admits coding (i.e., on singletons). For the result of this subsection, since T admitting coding on tuples immediately implies T is not monadically NIP, we could finish by [?BS, Theorem 8.1.8], which states that if T has IP then this is witnessed on singletons in a unary expansion. But the number of unary predicates used would depend on the length of the tuples in the tuple-coding configuration, which would weaken the results of Section 5.

Deriving non-structure results in a universal theory from the existence of a bad configuration is made much more involved if the configuration can occur on tuples. If one is willing to add unary predicates, arguments such as that from [?BS] mentioned above will often bring the configuration down to singletons. A general result in this case is [?Blum, Theorem 4.6] that (under mild assumptions) there is a formula defining the tuples of an indiscernible sequence in the expansion adding a unary predicate for each "coordinate strip" of the sequence. The results of [?Sim21] indicate the configuration can often be brought down to singletons just by adding parameters, instead of unary predicates, but these arguments seem difficult to adapt to tuple-coding configurations. Another approach, which we use here, is to take an instance of the configuration where the tuples have minimal length, and argue that the tuples then in many ways behave like singletons.

Definition 4.9. Given a tuple-coding configuration $\mathcal{I} = \langle \bar{a}_i : i \in I \rangle, \mathcal{J} = \langle \bar{b}_j : j \in J \rangle, \{ c_{i,j} \mid i \in I, j \in J \}, \phi(\bar{x}, \bar{y}, z)$, indexed by disjoint countable dense orderings $(I, \leq), (J \leq)$, an order-preserving permutation of $(I, \leq)$ (resp. $(J, \leq)$) naturally gives rise to permutation of $A = \bigcup \mathcal{I}$ (resp. $B = \bigcup \mathcal{J}$; call such permutations of A and B *standard permutations*.

A tuple-coding configuration as above is *regular* if $\mathfrak{C} \models \phi(\bar{d}, \bar{e}, c_{i,j}) \leftrightarrow \phi(\sigma(\bar{d}), \tau(\bar{e}), c_{i,j})$ whenever $\bar{d} \subseteq A, \bar{e} \subseteq B$ (including cases with $\bar{d} \notin \mathcal{I}, \bar{e} \notin \mathcal{J}$), σ is a standard permutation of A corresponding to an element of $Aut(I, \leq)$ fixing i, and τ is a standard permutation of B corresponding to an element of $Aut(J, \leq)$ fixing j.

By Ramsey and compactness, if T admits coding on tuples via the formula $\phi(\bar{x}, \bar{y}, z)$, then it admits a regular tuple-coding configuration via the same formula $\phi(\bar{x}, \bar{y}, z)$.

Definition 4.10. Let $\langle \bar{a}_i : i \in I \rangle, \langle \bar{b}_j : j \in J \rangle, \{ c_{i,j} \mid i \in I, j \in J \}, \phi(\bar{x}, \bar{y}, z)$ be a tuple-coding configuration with $(I, \leq), (J, \leq)$ countable, dense. The pair $(\bar{d}, \bar{e})$ with $\bar{d} \subseteq A, \bar{e} \subseteq B$ is a *witness* for $c_{k,\ell}$ if there are open intervals $I' \subseteq I, J' \subseteq J$ with $k \in I', \ell \in J'$ such that for all $k' \in I', \ell' \in J'$, we have $\mathfrak{C} \models \phi(\bar{d}, \bar{e}, c_{k',\ell'}) \iff (k, \ell) = (k', \ell')$.

A tuple-coding configuration *has unique witnesses up to permutation* if for every $c_{i,j} \in C$, the only witnesses for $c_{i,j}$ are of the form $(\sigma(\bar{a}_i), \tau(\bar{b}_j))$ for some σ a permutation of $\bar{a}_i$ and some τ a permutation of $\bar{b}_j$

Lemma 4.11. *Let $\langle \bar{a}_i : i \in I \rangle, \langle \bar{b}_j : j \in J \rangle, \{ c_{i,j} \mid i, j \in I \}, \phi(\bar{x}, \bar{y}, z)$ be a regular tuple-coding configuration for T, with $|\bar{x}| + |\bar{y}|$ minimal. Then this configuration has unique witnesses up to permutation.*

Proof. Suppose not, and let $(\bar{d}, \bar{e})$ be a witness for $c_{i,j}$, such that $(\bar{d}, \bar{e}) \neq (\sigma(\bar{a}_i), \tau(\bar{b}_j))$ for any σ, τ. First, if either $\bar{d} \cap \bar{a}_i = \emptyset$ or $\bar{e} \cap \bar{b}_j = \emptyset$, then regularity immediately implies that $(\bar{d}, \bar{e})$ is not a witness. So let $\bar{d}^*$ be the subsequence of $\bar{d}$ intersecting $\bar{a}_i$, and $\bar{e}^*$ the subsequence of $\bar{e}$ intersecting $\bar{b}_j$. Either $\bar{d}^* \neq \bar{d}$ or $\bar{e}^* \neq \bar{e}$ so assume the former.

Let $I^* \subseteq I$ be an open interval such that $\bar{d} \cap \bigcup(\bar{a}_i : i \in I^*) = \bar{d}^*$. Let $\phi^*(\bar{x}^*, \bar{y}, z)$ be the formula obtained by starting with $\phi(\bar{d}, \bar{y}, z)$, and then replacing the subtuple $\bar{d}^*$ with the variables $\bar{x}^*$; so we have plugged the elements of $\bar{d} \backslash \bar{d}^*$ as parameters into ϕ. For each $k \in I^*$, let $\bar{a}_k^*$ be the restriction of $\bar{a}_k$ to the coordinates corresponding to $\bar{d}^*$. Then $\langle \bar{a}_i^* : i \in I^* \rangle, \langle \bar{b}_j : j \in J \rangle, \{ c_{i,j} \mid i \in I^*, j \in J \}, \phi^*(\bar{x}^*, \bar{y}, z)$ is also a regular tuple-coding configuration, contradicting the minimality of $|\bar{x}| + |\bar{y}|$. □

The following Lemma completes the proof of Theorem 4.1.

Lemma 4.12. *Suppose T admits coding on tuples. Then T admits coding in an expansion by three unary predicates.*

Proof. Choose a tuple-coding configuration

$$\bar{A} = \langle \bar{a}_i : i \in I \rangle, \bar{B} = \langle \bar{b}_j : j \in J \rangle, C = \{ c_{i,j} \mid i \in I, j \in I \}, \phi(\bar{x}, \bar{y}, z)$$

with $|\bar{x}| + |\bar{y}|$ as small as possible. By the remarks following Definition 4.9. we may assume this configuration is regular, so by Lemma 4.11, it has unique witnesses up to permutation. Let $L^* = L \cup \{ A, B, C \}$ and let $\mathfrak{C}^*$ be the expansion of $\mathfrak{C}$ interpreting A as $\bigcup \bar{A}$, B as $\bigcup \bar{B}$, and C as itself. Let

$$\begin{aligned} &\phi^*(x, y, z) := A(x) \wedge B(y) \wedge C(z) \wedge \\ &\exists \bar{x}' \subseteq A, \bar{y}' \subseteq B(\phi(x\bar{x}', y\bar{y}', z) \wedge \forall z' \in C(z' \neq z \rightarrow \neg\phi(x\bar{x}', y\bar{y}', z'))) \end{aligned}$$

Let a_i be the first coordinate of $\bar{a}_i$, and b_j the first coordinate of $\bar{b}_j$. Then $A_1 = \{ a_i : i \in I \}, B_1 = \{ b_j : j \in J \}$, and C witness coding in $T^* = Th(\mathfrak{C}^*)$ via the L^*-formula $\phi^*(x, y, z)$. □

5. Finite structures

In this section, we restrict the language L of the theories we consider to be relational (i.e., no function symbols) with only finitely many constant symbols.

Definition 5.1. For a complete theory T and $M \models T$, $Age(M)$, the isomorphism types of finite substructures of M does not depend on the choice of M, so we let $Age(T)$ denote this class of isomorphism types.

The *growth rate* of $Age(T)$ (sometimes called the profile or (unlabeled) speed) is the function $\varphi_T(n)$ counting the number of isomorphism types with n elements in $Age(T)$.

We also investigate $Age(T)$ under the quasi-order of embeddability.We say $Age(T)$ is *well-quasi-ordered (wqo)* if this class does not contain an infinite antichain, and we say $Age(T)$ is *n-wqo* if $Age(M^*)$ is wqo for every expansion M^* of any model M of T by n unary predicates that partition the universe.

The definition of n-wqo is sometimes given for an arbitrary hereditary class $\mathcal{C}$ rather than an age, with $\mathcal{C}$ n-wqo if the class $\mathcal{C}^*$ containing every partition of every structure of $\mathcal{C}$ by at most n unary predicates remains wqo. Our definition is possibly weaker, but then its failure is stronger.

Example 1. Let $T = Th(\mathbb{Z}, succ)$. Then $Age(T)$ is wqo, but not 2-wqo, since $Age(T)$ contains arbitrarily long finite paths, and marking the endpoints of these paths with a unary predicate gives an infinite antichain.

By contrast, if $T = Th(\mathbb{Z}, \leq)$, then $Age(T)$ can be shown to be n-wqo for all n.

The following lemma shows that when considering n-wqo, adding finitely many parameters is no worse than adding another unary predicate.

Lemma 5.2. *Suppose $Age(M)$ is $(n+1)$-wqo. If M^* is an expansion of M by finitely many constants, then $Age(M^*)$ is n-wqo.*

Proof. Suppose an expansion by k constants is not n-wqo, as witnessed by an infinite antichain $\{ M_i^+ \}_{i\in\omega}$ in a language L^+ expanding the initial language by the k constants and by n unary predicates. Let M_i^* be the structure obtained from M_i^+ by forgetting the k constants, but naming their interpretations by a single new unary predicate. As $Age(M)$ is $(n+1)$-wqo, $\{ M_i^* \}_{i\in\omega}$ contains an infinite chain $M_{i_1}^* \hookrightarrow M_{i_2}^* \hookrightarrow \dots$ under embeddings. As there are only finitely many permutations of the constants, some embedding in the chain must preserve them, contradicting that $\{ M_i^+ \}_{i\in\omega}$ is an antichain. □

In both Theorem 5.3 and 5.6, the assumption that T has quantifier elimination is only used to get that the formula witnessing that T admits coding on tuples is quantifier-free, and the formula witnessing the order property in the stability part of Theorem 5.6, so the hypotheses of the theorems can be weakened to only these specific formulas being quantifier-free. This weakened assumption is used in [?ST]. From the proof of Proposition 3.11, if the failure of the f.s. dichotomy is witnessed by quantifier-free formulas, then the formula witnessing coding on tuples will be quantifier-free as well.

Theorem 5.3. *If a complete theory T has quantifier elimination in a relational language with finitely many constants is not monadically NIP, then $Age(T)$ has growth rate asymptotically greater than $(n/k)!$ for some $k \in \omega$ and is not 4-wqo.*

Proof. Since T is not monadically NIP, let

$$\langle \bar{a}_i : i \in I \rangle, \langle \bar{b}_j : j \in J \rangle, \{ c_{i,j} \mid i \in I, j \in J \}, \phi(\bar{x}, \bar{y}, z)$$

be a regular tuple-coding configuration with unique witnesses up to permutation. The only place we use T has QE is to choose ϕ quantifier-free. Let L^* expand by unary predicates for A, B, and C as well as constants for the parameters of ϕ, and let ϕ^* be as in the proof of 4.12. Let $\mathcal{A} \subseteq Age(T^*)$ be the set of finite substructures that can be constructed as follows.

(1) Pick $X \subset_{fin} I$, $Y \subset_{fin} J$, and $E \subset X \times Y$.

(2) Start with $\{\bar{a}_i : i \in X\} \cup \{\bar{b}_j : j \in Y\} \cup \{c_{i,j} \mid (i,j) \in E\}$.
(3) For every point $c_{i,j}$ added in the previous step, add the four elements $c_{i\pm\epsilon,j}$ and $c_{i,j\pm\epsilon}$, where $i \pm \epsilon$ are closer to i than any other element of X, and $j \pm \epsilon$ are closer to j than any other element of Y.
(4) Add the parameters of ϕ^*.

Claim. For any $M \in \mathcal{A}$ and $a, b, c \in M$, $\mathfrak{C} \models \phi^*(a,b,c) \iff M \models \phi^*(a,b,c)$.

Proof of Claim. Since ϕ is quantifier-free, it remains to check that if the existential quantifiers in ϕ^* are witnessed in $\mathfrak{C}$ and $\mathfrak{C} \models \phi^*(a,b,c)$ then they are witnessed in M, and if the universal fails in $\mathfrak{C}$ then it fails in M. From the unary predicates at the beginning of ϕ^*, we may let $a \in \bar{a}_i, b \in \bar{b}_j$, and $c = c_{k,\ell}$. If $\mathfrak{C} \models \phi(a,b,c)$, the only tuple in $\mathfrak{C}$ that can witness $\bar{x}'$ is the rest of the tuple $\bar{a}_i$, which will be in M because it only contains full tuples, and similarly for witnessing $\bar{y}'$. Since our configuration has unique witnesses up to permutation, if the universal quantifier fails in $\mathfrak{C}$, this is witnessed by an element $c_{k',\ell'}$ with $i - \epsilon \leq k' \leq i + \epsilon$ and $j - \epsilon \leq \ell' \leq j + \epsilon$. By regularity, this failure is also witnessed by some element in $\{c_{i\pm\epsilon,j}, c_{i,j\pm\epsilon}\}$. $\diamondsuit$

Given a bipartite graph G with n edges and no isolated vertices, we may encode it as a structure $M_G \in \mathcal{A}$ by starting with tuples $\bar{a}_i$ for each point in one part and tuples $\bar{b}_j$ for each point in the other part, and including $c_{k,\ell}$ whenever we want to encode an edge between $\bar{a}_k$ and $\bar{b}_\ell$. Note that $|M_G| = O(n)$, and this encoding preserves isomorphism in both directions. In the proof of [?Rap, Theorem 1.5], the asymptotic growth rate of such graphs is shown to be at least $(n/5)!$, which gives the desired growth rate for $Age(T^*)$ with the constant k depending on the length of the tuples in the tuple-coding configuration. Since expanding by finitely many unary predicates and constants increases the growth rate by at most an exponential factor, we also get the desired growth rate for $Age(T)$.

Furthermore, if M_H embeds into M_G, then H must be a (possibly non-induced) subgraph of G. So we get that $Age(T^*)$ is not wqo by encoding even cycles. We expanded by three unary predicates, and by Lemma 5.2 the parameters may be replaced by another unary predicate while still preserving the failure of wqo, so we get that $Age(T)$ is not 4-wqo. $\square$

Remark 5.4. There is a homogeneous structure, with automorphism group $S_\infty \operatorname{Wr} S_2$ in its product action, that is not monadically NIP and whose growth rate is the number of bipartite graphs with a prescribed bipartition, n edges, and no isolated vertices. So the lower bound in this theorem cannot be raised above the growth rate of such graphs. Precise asymptotics for this growth rate are not known, although it is slower than $n!$ and [?CPS, Theorem 7.1] improves Macpherson's lower bound to $(\frac{n}{\log n^{2+\epsilon}})^n$ for every $\epsilon > 0$.

If Conjecture 1 from the Introduction (in particular (1) $\Rightarrow$ (2)) is confirmed, then the lower bound on the growth rate in Theorem 5.3 would also confirm [?Rap, Conjecture 3.5] that for homogeneous structures there is a gap from exponential growth rate to growth rate at least $(n/k)!$ for some $k \in \omega$.

Theorem 5.3 is somewhat surprising. Since passing to substructures can be simulated by adding unary predicates, it is clear that if T is monadically tame, then $Age(T)$ should be tame. However, unary predicates can do more, so it seems plausible that $Age(T)$ could be tame even though T is not monadically tame. Our next theorem gives some explanation for why this does not occur, at least when assuming quantifier elimination.

First we need to define stability and NIP for hereditary classes. The following definition is standard and appears, for example, in [?ST, §8.1].

Definition 5.5. For a formula $\phi(\bar{x}, \bar{y})$ and a bipartite graph $G = (I, J, E)$, we say a structure M *encodes* G *via* ϕ if there are sets $A = \{ \bar{a}_i \mid i \in I \} \subseteq M^{|x|}, B = \{ \bar{b}_j \mid j \in J \} \subseteq M^{|y|}$ such that $M \models \phi(\bar{a}_i, \bar{b}_j) \Leftrightarrow G \models E(i, j)$.

A class of structures $\mathcal{C}$ has *IP* if there is some formula $\phi(\bar{x}, \bar{y})$ such that for every finite, bipartite graph $G = (I, J, E)$, there is some $M_G \in \mathcal{C}$ encoding G via ϕ. Otherwise, $\mathcal{C}$ is *NIP*.

A class of structures $\mathcal{C}$ is *unstable* if there is some formula $\phi(\bar{x}, \bar{y})$ such that for every finite half-graph G, there is some $M_G \in \mathcal{C}$ encoding G via ϕ. Otherwise, $\mathcal{C}$ is *stable*.

Equivalently, by compactness arguments, $\mathcal{C}$ is NIP (resp. stable) if and only if every completion of $Th(\mathcal{C})$, the common theory of structures in $\mathcal{C}$, is. Note that it suffices to witness that $\mathcal{C}$ has IP or is unstable using a formula with parameters, since we can remove them by appending the parameters to each $\bar{a}_i$.

The sort of collapse between monadic NIP and NIP in hereditary classes observed in the next theorem occurs for binary ordered structures [?ST], since there the formula giving coding on tuples is quantifier-free. It also occurs for monotone graph classes (i.e. specified by forbidding non-induced subgraphs), where NIP actually collapses to monadic stability, and agrees with nowhere-denseness [?AA].

Theorem 5.6. *Suppose that a complete theory T in a relational language with finitely many constants has quantifier elimination. Then $Age(T)$ is NIP if and only if T is monadically NIP, and $Age(T)$ is stable if and only if T is monadically stable.*

Proof. We first consider the NIP case.

($\Leftarrow$) Suppose $Age(T)$ has IP, as witnessed by the formula $\phi(\bar{x}, \bar{y})$. By compactness, there is a model N of the universal theory of T in which ϕ encodes the generic bipartite graph. But then N is a substructure of some $M \models T$, and naming a copy of N in M by a unary predicate U and relativizing ϕ to U gives a unary expansion of M with IP.

($\Rightarrow$) Suppose T is not monadically NIP, witnessed by a tuple-coding configuration $\mathcal{I} = (\bar{a}_i : i \in I), \mathcal{J} = (\bar{b}_j : j \in J), C = \{ c_{i,j} \mid i \in I, j \in J \}, \phi(\bar{x}, \bar{y}, z)$, with ϕ quantifier-free and containing parameters $\bar{m}$. By Remark 3.12, we may also assume the configuration is tidy. For any bipartite graph G, let $M_G \in Age(T)$ contain $\bar{m}$, tuples from $\mathcal{I}$ and $\mathcal{J}$ corresponding to the two parts of G, and an element of $c_{i,j}$ for each edge of G so that $R^*(\bar{x}, \bar{y}; \bar{m}) := \exists z \in C(\phi(\bar{x}, \bar{y}, z; \bar{m}))$ encodes G on $\mathcal{I}(M_G) \times \mathcal{J}(M_G)$. But by tidiness, $R(\bar{x}, \bar{y}; \bar{m}) := \exists z(\phi(\bar{x}, \bar{y}, z; \bar{m}) \wedge z \notin \bar{m})$ encodes G on $\mathcal{I}(M_G) \times \mathcal{J}(M_G)$ as well.

For the stable case, the backwards direction is the same except using the infinite half-graph in place of the generic bipartite graph. For the forwards direction, if T is unstable then by quantifier-elimination $Age(T)$ is also unstable. If T is stable but not monadically stable, then by [?BS, Lemma 4.2.6] T is not monadically NIP, so we are finished by the NIP case. $\square$

References