

Algorithmic Model Theory

SS 2016

Prof. Dr. Erich Grädel and Dr. Wied Pakusa

Mathematische Grundlagen der Informatik
RWTH Aachen

Contents

1	The classical decision problem	1
1.1	Basic notions on decidability	2
1.2	Trakhtenbrot's Theorem	7
1.3	Domino problems	14
1.4	Applications of the domino method	17
1.5	The finite model property	20
1.6	The two-variable fragment of FO	22
2	Descriptive Complexity	31
2.1	Logics Capturing Complexity Classes	31
2.2	Fagin's Theorem	33
2.3	Second Order Horn Logic on Ordered Structures	38
3	Expressive Power of First-Order Logic	43
3.1	Ehrenfeucht-Fraïssé Theorem	43
3.2	Hanf's technique	47
3.3	Gaifman's Theorem	49
3.4	Lower bound for the size of local sentences	54
4	Zero-one laws	61
4.1	Random graphs	61
4.2	Zero-one law for first-order logic	63
4.3	Generalised zero-one laws	67
5	Modal, Inflationary and Partial Fixed Points	73
5.1	The Modal μ -Calculus	73
5.2	Inflationary Fixed-Point Logic	75
5.3	Simultaneous Inductions	81
5.4	Partial Fixed-Point Logic	82



This work is licensed under:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de/>

Dieses Werk ist lizenziert unter:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de/>

© 2016 Mathematische Grundlagen der Informatik, RWTH Aachen.

<http://www.logic.rwth-aachen.de>

4 Zero-one laws

4.1 Random graphs

We consider the class \mathcal{G}_n of (undirected) graphs over $\{0, \dots, n - 1\}$, i.e.

$$\mathcal{G}_n := \{G = (V, E) : G \text{ graph}, V = \{0, \dots, n - 1\}\},$$

In order to introduce *random graphs* we consider a sequence of probability distributions $\bar{\mu} = (\mu_1, \mu_2, \dots)$ on $(\mathcal{G}_1, \mathcal{G}_2, \dots)$, i.e. $\mu_n : \mathcal{G}_n \rightarrow [0, 1]$ and $\sum_{G \in \mathcal{G}_n} \mu(G) = 1$ for all $n \geq 1$. This defines a sequence of probability spaces $(\mathcal{G}_1, \mu_1), (\mathcal{G}_2, \mu_2), \dots$ on classes of graphs of increasing size.

Example 4.1.

(1) The *uniform distribution* μ_n assigns equal probability to each graph:

$$\mu_n(G) = \frac{1}{2^{\binom{n}{2}}}.$$

(2) Let $p : \mathbb{N} \rightarrow [0, 1]$ be an arbitrary mapping. Then the probability space $\mathcal{G}_{n,p} = (\mathcal{G}_n, \mu_{p,n})$ is defined by the following random experiment: determine for every pair (u, v) with $0 \leq u < v < n$ whether $(u, v) \in E$ using a random variable X taking values 0, 1 (False and True) with $\Pr[X = 1] = p(n)$ and $\Pr[X = 0] = (1 - p(n))$. Observe that for $p = \frac{1}{2}$ one obtains the uniform distribution.

We make the following convention: unless otherwise stated, μ_n denotes the uniform distribution. For a class \mathcal{K} of graphs we set

$$\mu_n(\mathcal{K}) := \mu_n(\mathcal{K} \cap \mathcal{G}_n) = \sum_{G \in \mathcal{K} \cap \mathcal{G}_n} \mu_n(G).$$

This definition formalises what it means that a random graph $G \in \mathcal{G}_n$ has a certain property \mathcal{K} . However, in what follows, we are not interested

in random graphs of some fixed size $n \in \mathbb{N}$ but much more in the behaviour of the probability $\mu_n(K)$ if we increase the size of graphs, i.e. if we let n approach infinity.

Definition 4.2. The *asymptotic probability* of a class \mathcal{K} of graphs (with respect to $\bar{\mu}$) is defined as

$$\mu(\mathcal{K}) := \lim_{n \rightarrow \infty} \mu_n(\mathcal{K}),$$

in the case that this sequence has a limit. In particular, if ψ is a sentence over vocabulary $\{E\}$ in some logic \mathcal{L} , then the *asymptotic probability* of ψ (with respect to $\bar{\mu}$) is defined as

$$\mu(\psi) := \lim_{n \rightarrow \infty} \mu_n(\{G \in \mathcal{G}_n : G \models \psi\}),$$

again only for the case that the limit exists.

Example 4.3.

(1) Let $\mathcal{K} = \{G : G \text{ is a clique}\}$. Then

$$\lim_{n \rightarrow \infty} \mu_n(\mathcal{K}) = \lim_{n \rightarrow \infty} \frac{1}{2^{\binom{n}{2}}} = 0.$$

(2) Let H be a graph and let $\mathcal{K}_H = \{G : G \text{ contains } H \text{ as subgraph}\}$.

For $n > k \cdot |H|$ we have

$$\mu_n(\mathcal{K}_H) \geq 1 - (1 - (2^{-|E(H)|}))^k,$$

hence $\mu(\mathcal{K}_H) = 1$ since $k \rightarrow \infty$ for $n \rightarrow \infty$.

(3) Let $\mathcal{K} = \{G : G \text{ is three-colourable}\}$. Then

$$\lim_{n \rightarrow \infty} \mu_n(\mathcal{K}) \leq 1 - \lim_{n \rightarrow \infty} \mu_n(\{G \in \mathcal{G}_n : G \text{ contains } K_4\}) = 0.$$

(4) Recall that we have $\lim_{n \rightarrow \infty} \mu_n(\{G : (3, 17) \in E\}) = \frac{1}{2}$.

(5) The asymptotic probability is not defined for every class of graphs.

For instance, consider $\mathcal{K} = \{G : G \text{ has an even number of nodes}\}$.

Then the sequence $(\mu_n(\mathcal{K}))_{n \geq 1} = (0, 1, 0, 1, \dots)$ has no limit.

4.2 Zero-one law for first-order logic

In this section we prove the *zero-one law* for first-order logic:

Theorem 4.4. For sentences $\psi \in \text{FO}$ (over relational vocabulary) we have

$$\mu(\psi) = 0 \quad \text{or} \quad \mu(\psi) = 1.$$

To put it in words, every first-order definable property of graphs either holds *almost never* or *almost surely* on random graphs of increasing size.

Definition 4.5. An *atomic graph k -type* is a maximal consistent set t of $\text{FO}(\{E\})$ -literals in variables x_1, \dots, x_k , i.e. $Ex_i x_j, \neg Ex_i x_j, x_i = x_j, x_i \neq x_j$, which is consistent with the graph axioms $(\forall x_1 \forall x_2 (\neg Ex_1 x_1 \wedge Ex_1 x_2 \leftrightarrow Ex_2 x_1))$. Furthermore, for a graph $G = (V, E)$ and $\bar{a} \in V^k$ we define the *atomic graph k -type of \bar{a}* by

$$t_G(\bar{a}) := \{\varphi(x_i, x_j) : \varphi \text{ an FO}(\{E\})\text{-literal such that } G \models \varphi(a_i, a_j)\}.$$

Formally, an atomic k -type t is a set but we frequently identify it with the formula $t(\bar{x}) = \bigwedge_{\varphi \in t} \varphi(\bar{x})$ (this formula is an FO-formula, since there are only finitely many $\{E\}$ -literals in k variables).

In what follows, let $s(\bar{x})$ and $t(\bar{x})$ denote atomic graph types of tuples of distinct elements, i.e. $s, t \models \bigwedge_{i < j \leq k} x_i \neq x_j$. We say that an atomic $(m+1)$ -type $t(x_1, \dots, x_m, x_{m+1})$ *extends* an atomic m -type $s(x_1, \dots, x_m)$ if $s \subseteq t$, or equivalently, if $t \models s$.

Definition 4.6. Let $s(x_1, \dots, x_m)$ and $t(x_1, \dots, x_m, x_{m+1})$ be atomic types such that $s \subseteq t$. We define the *extension axiom* $\sigma_{s,t}$ by

$$\sigma_{s,t} := \forall x_1 \dots \forall x_m (s(\bar{x}) \rightarrow \exists x_{m+1} t(\bar{x}, x_{m+1})).$$

Furthermore, we let T be the set of all extension axioms together with the graph axioms.

The proof of the zero-one law for FO relies on the following properties of the extension axioms and the set T :

(1) $\mu(\sigma_{s,t}) = 1$ for all $\sigma_{s,t} \in T$.

- (2) T is ω -categorical, i.e. there is, up to isomorphism, only one countable model of T . This structure is known as the *Rado graph*.
- (3) T is complete, i.e. for all $\psi \in \text{FO}$ either $T \models \psi$ or $T \models \neg\psi$.

We proceed to establish these three properties.

Lemma 4.7. Let $\sigma_{s,t} \in T$ be an extension axiom. Then $\mu(\sigma_{s,t}) = 1$.

Proof. Let $\sigma_{s,t} := \forall x_1 \cdots \forall x_m (s(\bar{x}) \rightarrow \exists x_{m+1} t(\bar{x}, x_{m+1}))$. For every $i = 1, \dots, m$ we have $t \models Ex_i x_{m+1}$ or $t \models \neg Ex_i x_{m+1}$. Let $G \in \mathcal{G}_n$ be a random graph and $a_1, \dots, a_m \in \{0, \dots, n-1\}$. For every fixed $a_{m+1} \in V \setminus \{a_1, \dots, a_m\}$, the experiments $G \models Ea_i a_{m+1}$ are stochastically independent and have probability $\frac{1}{2}$. Hence

$$\Pr[G \models t(\bar{a}, a_{m+1}) | G \models s(\bar{a})] = \frac{1}{2^m}.$$

Thus, probability that *no* element $a_{m+1} \in V \setminus \{a_1, \dots, a_m\}$ extends a realisation \bar{a} of s to a realisation of (\bar{a}, a_{m+1}) of t is $(1 - \frac{1}{2^m})^{n-m}$. In conclusion, we obtain

$$\begin{aligned} \mu_n(\neg\sigma_{s,t}) &= \mu_n(\exists x_1 \cdots \exists x_n (s(\bar{x}) \wedge \forall x_{m+1} \neg t(\bar{x}, x_{m+1}))) \\ &\leq n^m \cdot (1 - \frac{1}{2^m})^{n-m} \xrightarrow{\text{exp. fast}} 0, \end{aligned}$$

and thus $\mu(\sigma_{s,t}) = 1$. Q.E.D.

The compactness theorem implies that also every logical consequence of the extensions axioms almost surely holds in a random graph.

Corollary 4.8. If $T \models \psi$ then $\mu(\psi) = 1$, and the set T is satisfiable.

Proof. If $T \models \psi$, then by the compactness theorem there is a finite set $T_0 \subseteq T$ such that $T_0 \models \psi$. Hence, we have $\mu_n(\psi) \geq \prod_{\sigma \in T_0} \mu_n(\sigma)$ for all $n \geq 1$ and thus $\lim_{n \rightarrow \infty} \mu_n(\psi) = 1$ by Lemma 4.7. In particular $T \models \forall x(x \neq x)$ since $\mu(\forall x(x \neq x)) = 0$. Q.E.D.

Interestingly, one can give explicit description of models of T and we present two different possibilities here. However, as we show later that T is ω -categorical, these models are isomorphic.

Definition 4.9 (Rado graph). The following graphs are models of T .

- (1) Let p_i denote the i -th prime number. We define $G = (\mathbb{N}, E)$ with

$$E := \{(i, j) \in \mathbb{N} \times \mathbb{N} : p_i \mid j \text{ or } p_j \mid i.\}$$

We claim that $G \models T$. To see this, we choose an arbitrary extension axiom $\sigma_{s,t} := \forall x_1 \cdots \forall x_m (s(\bar{x}) \rightarrow \exists x_{m+1} t(\bar{x}, x_{m+1})) \in T$.

Let $I \cup J = \{1, \dots, m\}$ be the partition defined by t with respect to the following condition

- If $t \models Ex_i x_{m+1}$ then $i \in I$, and
- if $t \models \neg Ex_i x_{m+1}$ then $i \in J$.

Let $a_1, \dots, a_k \in A$ such that $G \models s(a_1, \dots, a_k)$. We set $a_{m+1} := \prod_{i \in I} p_{a_i} q$ where q is a prime number with $q > p_{a_1} \cdots p_{a_m}$. Then it is easy to check that $G \models Ea_i a_{m+1}$ for all $i \in I$ and $G \models \neg Ea_j a_{m+1}$ for all $j \in J$.

- (2) The set HF of *hereditarily finite sets* is defined by:

- $\emptyset \in \text{HF}$
- If $a_1, \dots, a_k \in \text{HF}$, then also $\{a_1, \dots, a_k\} \in \text{HF}$.

Let $G = (\text{HF}, E)$ with $E := \{(a, b) : a \in b \text{ or } b \in a\}$. Similarly as above, one can show that $G \models T$.

Theorem 4.10. Let $G = (V_G, E_G)$ and $H = (V_H, E_H)$ be two countable models of T . Then $G \cong H$. The unique countable model of T is known as the *Rado graph* \mathcal{R} .

Proof. First of all, it is clear that T has no finite models, hence G and H are infinite graphs. We fix two enumerations of V_G and V_H and inductively construct a sequence of partial isomorphism p_0, p_1, p_2, \dots between G and H such that $p_0 \subseteq p_1 \subseteq p_2 \subseteq \dots$. For the base case, we set $p_0 := \emptyset$. For the induction step let $p_i = \{(a_1, b_1), \dots, (a_i, b_i)\} \in \text{Loc}(G, H)$ be already defined. We distinguish between the following two cases:

- If i is even, choose $a_{i+1} \in V_G$ to be the minimal element (with respect to the enumeration of V_G) which is not in the domain of p_i , i.e. $a_{i+1} \notin \{a_1, \dots, a_i\}$. Let $s := t_G(a_1, \dots, a_i)$ and $t := t_G(a_1, \dots, a_{i+1})$. Since p_i is a partial isomorphism we know that $H \models s(b_1, \dots, b_i)$.

Since $H \models \sigma_{s,t}$ there exists an element $b_{i+1} \in V_H$ such that $H \models t(b_1, \dots, b_{i+1})$. We set $p_{i+1} := p_i \cup \{(a_{i+1}, b_{i+1})\}$ and obtain a partial isomorphism extending p_i .

- If i is odd, we proceed analogously, but this time we let $b_{i+1} \in V_H$ be the minimal element (with respect to the enumeration of V_H) which is not in the image of p_i , i.e. $b_{i+1} \notin \{b_1, \dots, b_i\}$. For $s := t_H(b_1, \dots, b_i)$ and $t := t_H(b_1, \dots, b_{i+1})$, the same reasoning as above yields an element $a_{i+1} \in V_G$ such that $G \models t(a_1, \dots, a_{i+1})$. Again we obtain an extended partial isomorphism by setting $p_{i+1} := p_i \cup \{(a_{i+1}, b_{i+1})\}$.

Finally we let $p := \bigcup_{i \geq 0} p_i$. By construction we have that $\text{dom}(p) = V_G$ and $\text{im}(p) = V_H$, hence $p : G \xrightarrow{\sim} H$. Q.E.D.

In particular, ω -categorical theories are complete:

Theorem 4.11. T axiomatises a complete theory, i.e. for all sentences $\psi \in \text{FO}(\{E\})$ we have $T \models \psi$ or $T \models \neg\psi$.

Proof. Assume for some sentence $\psi \in \text{FO}(\{E\})$ it holds that $T \not\models \psi$ and $T \not\models \neg\psi$. Then by the downwards Löwenheim-Skolem theorem, there exist two countable graphs G and H with $G \models T \cup \{\psi\}$ and $H \models T \cup \{\neg\psi\}$. In particular this implies $G \not\cong H$, which contradicts Theorem 4.10. Q.E.D.

Theorem 4.12. [Glebskiĭ et al., R. Fagin] For all $\psi \in \text{FO}(\{E\})$ it holds:

$$\mu(\psi) = 0 \quad \text{or} \quad \mu(\psi) = 1.$$

Proof. If $T \models \psi$, then $\mu(\psi) = 1$. Otherwise, $T \models \neg\psi$, and hence $\mu(\psi) = 1 - \mu(\neg\psi) = 0$. Q.E.D.

In particular, we can give a precise characterisation of those first-order properties which hold almost surely in random graphs.

Corollary 4.13. Let $\psi \in \text{FO}(\{E\})$. Then

$$\mu(\psi) = 1 \quad \text{iff} \quad T \models \psi \quad \text{iff} \quad \mathcal{R} \models \psi.$$

4.2.1 Applications

We can use Theorem 4.12 to show that certain classes of graphs are not definable in first-order logic: if a class \mathcal{K} of graphs has undefined asymptotic probability or an asymptotic probability different from 0 and 1, then clearly \mathcal{K} cannot be defined in first-order logic. More generally, this method yields non-definability of \mathcal{K} for *every* logic that has a 0-1-law, e.g. for $L_{\infty\omega}^\omega$ as we see later. For instance, consider the class $\text{EvenV} = \{G = (V, E) : |V| \text{ is even}\}$ with undefined asymptotic probability or the class $\text{EvenE} = \{G = (V, E) : |E| \text{ is even}\}$ with $\mu(\text{EvenE}) = \frac{1}{2}$. Moreover, we can use our results as a convenient method to determine the asymptotic probability for many natural classes of graphs.

- (1) We want to determine $\mu(\text{Con})$ where Con denotes the class of connected graphs. Let s be an atomic 2-type in variables x, y containing $\neg Exy$ and let t be the atomic 3-type in variables x, y, z which extends s and contains $Exz \wedge Eyz$. Then $G \models \sigma_{s,t}$ iff G has diameter at most 2. Hence, $G \models \sigma_{s,t}$ implies $G \in \text{Con}$, which means that $\mu(\text{Con}) = 1$.
- (2) Let \mathcal{K} be any class of graphs which exclude a forbidden subgraph $H = (\{v_1, \dots, v_k\}, E)$. Then $\mu(\mathcal{K}) = 0$. To see this, we set $s_i(x_1, \dots, x_i) := t_H(v_1, \dots, v_i)$ for $i \leq k$ and consider the extension axioms $\sigma_{s_i, s_{i+1}}$. Then clearly $\psi := \bigwedge_{i < k} \sigma_{s_i, s_{i+1}}$ is a logical consequence of T , which means that $\mu(\psi) = 1$. Moreover, if $G \models \psi$, then G contains H as an induced subgraph. We conclude that $\mu(\mathcal{K}) \leq 1 - \mu(\psi) = 0$. As an application, consider the class of planar graphs which exclude K_5 (the complete graph on 5 vertices) and the class of k -colourable graphs which exclude K_{k+1} (where k is fixed). To put it in words, a random graph is almost never planar nor k -colourable.

4.3 Generalised zero-one laws

In this section we want to generalise our considerations in two different ways. Firstly, instead of restricting ourselves to graphs, we want to work on more general classes of structures and analyse whether the zero-one-

law for FO still holds. Secondly, as FO has rather limited expressive power, we look for zero-one laws for more powerful logics as well.

Let τ be an arbitrary vocabulary (not necessarily relational). By $\text{Str}_n(\tau)$ we denote the set of all τ -structures over the universe $\{0, \dots, n-1\}$. As before we define a sequence $\bar{\mu} = (\mu_1, \mu_2, \dots)$ of uniform probability distributions $\mu_n : \text{Str}_n(\tau) \rightarrow [0, 1]$, i.e. for every $\mathfrak{A} \in \text{Str}_n(\tau)$ we set

$$\mu_n(\mathfrak{A}) = \frac{1}{|\text{Str}_n(\tau)|}.$$

We claim that $\text{FO}(\tau)$ has a zero-one law if, and only if, τ contains no function symbols. To this end, we first consider the case where τ contains function symbols:

- (1) Assume $\{P, c\} \subseteq \tau$ where c is a constant symbol and P a monadic relation. Then for $\psi := Pc$ we have $\mu_n(\psi) = \frac{1}{2}$ for all $n \geq 1$, hence $\mu(\psi) = \frac{1}{2}$, i.e. the zero-one law does not hold in this case.
- (2) Assume $f \in \tau$ where f is a unary function symbol. Consider the $\text{FO}(\tau)$ -sentence $\psi := \exists x(fx = x)$ stating that f has a fixed point. For $n \geq 1$ we have

$$\mu_n(\psi) = 1 - \prod_{i=0}^{n-1} \underbrace{\left(\frac{n-1}{n}\right)}_{=\Pr[f(i) \neq i]} = 1 - \left(1 - \frac{1}{n}\right)^n.$$

Since $\left(1 - \frac{1}{n}\right)^n \rightarrow e^{-1}$ for $n \rightarrow \infty$, the zero-one law does not hold in this case either.

For the other direction, let τ be purely relational, $\tau = \{R_1, \dots, R_k\}$. The proof strategy we used over graphs generalises for this general in a straightforward way:

- An *atomic τ -type in k variables* is a maximal, consistent set of τ -literals over variables x_1, \dots, x_k . For a τ -structure \mathfrak{A} and $\bar{a} \in \mathfrak{A}$ we set $t_{\mathfrak{A}}(\bar{a}) = \{\varphi(\bar{x}) : \varphi \text{ a } \tau\text{-literal with } \mathfrak{A} \models \varphi(\bar{a})\}$.
- The *τ -extension axiom* $\sigma_{s,t}$ for two atomic τ -types s and t (in k and

$k+1$ variables, respectively) with $s \subseteq t$ is defined as

$$\sigma_{s,t} := \forall \bar{x}(s(\bar{x}) \rightarrow \exists x_{k+1}t(\bar{x}, x_{k+1})).$$

As before, we let T denote the set of all τ -extension axioms

- Again we can show that $\mu(\sigma_{s,t}) = 1$ for all $\sigma_{s,t} \in T$. Let r denote the number of literals in t which contain x_{m+1} . Then, for a random structure $\mathfrak{A} \in \text{Str}_n(\tau)$, $\bar{a} \in A$ and a_{m+1} it holds

$$\Pr[\mathfrak{A} \models t(\bar{a}, a_{m+1}) \mid \mathfrak{A} \models s(\bar{a})] = 2^{-r}.$$

Thus

$$\begin{aligned} \mu_n(\neg\sigma_{s,t}) &= \mu_n(\exists \bar{x}(s(\bar{x}) \wedge \forall x_{m+1} \neg t(\bar{x}, x_{m+1}))) \\ &\leq n^m (1 - 2^{-r})^{n-m} \xrightarrow{\text{exp. fast}} 0. \end{aligned}$$

- T is ω -categorical: analogously!

Our analysis raises the question why even basic functions but not arbitrary relations inhibit a zero-one law. The reason is that atomic experiments are not longer stochastically independent. For instance, consider the experiments $f(a) = b$ and $f(a) = c$ (for $b \neq c$), then $\Pr[f(a) = c \mid f(a) = b] = 0 \neq \Pr[f(a) = c]$.

4.3.1 Zero-one law for $L_{\infty\omega}^\omega$

We proceed to show that the zero-one law holds for $L_{\infty\omega}^\omega$ as well (restricted to relational vocabularies). In particular, since $\text{LFP} \leq L_{\infty\omega}^\omega$, this means that a random graph either almost surely has an LFP-definable property or almost never does. With FO^k we denote the k -variable fragment of FO, i.e. $\text{FO}^k = \text{FO} \cap L_{\infty\omega}^k = \{\varphi \in \text{FO} : \varphi \text{ only contains variables } x_1, \dots, x_k\}$. If we restrict the set of extension axioms T to FO^k we obtain finite sets of approximations of T which are again sentences in FO^k ; more specifically, we set

$$\Theta_k := \bigwedge T \cap \text{FO}^k = \bigwedge \{\sigma_{s,t} : \sigma_{s,t} \in T \cap \text{FO}^k\} \in \text{FO}^k.$$

The central property of these approximations for T is stated in the following theorem: in models of Θ_k , every $L_{\infty\omega}^k$ -formula is equivalent to a simple Boolean combinations of atomic k -types. In particular, every $L_{\infty\omega}^k$ -sentence is either true or false in all models of Θ_k .

Theorem 4.14. Let $m \leq k$, $s(x_1, \dots, x_m)$ an atomic m -type and $\varphi(x_1, \dots, x_m) \in L_{\infty\omega}^k$. Then

$$\begin{array}{ll} \text{either} & \Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \varphi(\bar{x})) \\ \text{or} & \Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \neg\varphi(\bar{x})). \end{array}$$

Proof. We proceed by induction on φ and simultaneously show the claim for all $m \leq k$ and atomic types s . If φ is atomic, then either $\varphi \in s$ or $\neg\varphi \in s$. If $\varphi = \neg\psi$, the claim directly follows.

Let $\varphi = \bigwedge \Psi$, $\Psi \subseteq L_{\infty\omega}^k$. By induction hypothesis for all $\psi \in \Psi$

$$\begin{array}{ll} \text{either} & \Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \psi(\bar{x})) \\ \text{or} & \Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \neg\psi(\bar{x})). \end{array}$$

If $\Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \psi(\bar{x}))$ for all $\psi \in \Psi$, then $\Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \bigwedge \Psi(\bar{x}))$. Otherwise, $\Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \neg \bigwedge \Psi(\bar{x}))$.

Let $\varphi(\bar{x}) = \exists y\psi(\bar{x}, y)$ and assume that $\Theta_k \not\models \forall \bar{x}(s(\bar{x}) \rightarrow \neg\varphi(\bar{x}))$. Choose a structure $\mathfrak{A} \models \Theta_k$ with $\mathfrak{A} \models \exists \bar{x}(s(\bar{x}) \wedge \exists y\psi(\bar{x}, y))$ and consider the following two cases

- If $y \notin \{x_1, \dots, x_m\}$, i.e. $y \in \{x_{m+1}, \dots, x_k\}$; let $a_1, \dots, a_m, b \in A$ such that $\mathfrak{A} \models s(\bar{a}) \wedge \psi(\bar{a}, b)$. We define the atomic type $t(x_1, \dots, x_m, y) := t_{\mathfrak{A}}(\bar{a}, b)$ with $s \subseteq t$. In particular,

$$\mathfrak{A} \models \exists \bar{x}\exists y(t(\bar{x}, y) \wedge \psi(\bar{x}, y)).$$

By induction hypothesis we know that

$$\mathfrak{A} \models \forall \bar{x}\forall y(t(\bar{x}, y) \rightarrow \psi(\bar{x}, y)),$$

and since $\sigma_{s,t} = \forall \bar{x}(s(\bar{x}) \rightarrow \exists y t(\bar{x}, y))$ is an extension axiom con-

tained in Θ_k we finally obtain

$$\mathfrak{A} \models \forall \bar{x}(s(\bar{x}) \rightarrow \exists y\psi(\bar{x}, y)).$$

- If $y \in \{x_1, \dots, x_m\}$, i.e. $y = x_j$ for $j \leq m$; let $\bar{a} \in A$ such that $\mathfrak{A} \models s(\bar{a}) \wedge \exists x_j\psi(\bar{a})$, and let \bar{x}^* and \bar{a}^* denote the tuples \bar{x} and \bar{a} without the j -th component, i.e.

$$\begin{aligned} \bar{x}^* &:= x_1 \cdots x_{j-1} x_{j+1} \cdots x_k \\ \bar{a}^* &:= a_1 \cdots a_{j-1} a_{j+1} \cdots a_m. \end{aligned}$$

Similarly, let $s^*(\bar{x}^*) := t_{\mathfrak{A}}(\bar{a}^*)$ be the atomic type of \bar{a}^* in \mathfrak{A} . Then $s^* \subseteq s$ and there is $b \in A$ such that

$$\mathfrak{A} \models s^*(\bar{a}^*) \wedge \psi(\bar{a} \frac{b}{a_j}), \quad \text{where } \bar{a} \frac{b}{a_j} := a_1 \cdots a_{j-1} b a_{j+1} \cdots a_m.$$

For $t^*(\bar{x}) := t_{\mathfrak{A}}(\bar{a} \frac{b}{a_j})$ we thus have $\mathfrak{A} \models \exists(t^*(\bar{x}) \wedge \psi(\bar{x}))$, and the induction hypothesis yields

$$\Theta_k \models \forall \bar{x}(t^*(\bar{x}) \rightarrow \psi(\bar{x})).$$

As above, since $s^* \subseteq t^*$, it holds that $\Theta_k \models \forall \bar{x}^*(s^*(\bar{x}^*) \rightarrow \exists x_j t^*(\bar{x}))$, and altogether we obtain

$$\Theta_k \models \forall \bar{x}(s(\bar{x}) \rightarrow \exists x_j \psi(\bar{x})).$$

Q.E.D.

Corollary 4.15. For every $L_{\infty\omega}^k$ -sentence ψ we either have $\Theta_k \models \psi$ or $\Theta_k \models \neg\psi$.

Corollary 4.16. If $\mathfrak{A} \models \Theta_k$ and $\mathfrak{B} \models \Theta_k$, then $\mathfrak{A} \equiv_{L_{\infty\omega}^k} \mathfrak{B}$.

Corollary 4.17 (Kolaitis, Varidi 1992). For every sentence $\psi \in L_{\infty\omega}^\omega$ (over a relational signature) we have $\mu(\psi) = 0$ or $\mu(\psi) = 1$.

Proof. Let $\psi \in L_{\infty\omega}^k$ for some $k \geq 1$. Then by Corollary 4.15 we have $\Theta_k \models \psi$ or $\Theta_k \models \neg\psi$. Since $\Theta_k \subseteq T$ is finite, we have $\mu(\Theta_k) = 1$ and thus the claim follows. Q.E.D.