THE BULLETIN OF SYMBOLIC LOGIC
Volume 2, Number 4, Dec. 1996

ALMOST EVERYWHERE EQUIVALENCE OF LOGICS IN FINITE MODEL THEORY

LAURI HELLA, PHOKION G. KOLAITIS, AND KERKKO LUOSTO

Abstract. We introduce a new framework for classifying logics on finite structures and studying their expressive power. This framework is based on the concept of *almost everywhere equivalence of logics*, that is to say, two logics having the same expressive power on a class of asymptotic measure 1. More precisely, if $\mathcal{L}$, $\mathcal{L}'$ are two logics and μ is an asymptotic measure on finite structures, then $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu)$ means that there is a class $\mathcal{C}$ of finite structures with $\mu(C) = 1$ and such that $\mathcal{L}$ and $\mathcal{L}'$ define the same queries on $\mathcal{C}$. We carry out a systematic investigation of $\equiv_{\text{a.e.}}$ with respect to the uniform measure and analyze the $\equiv_{\text{a.e.}}$-equivalence classes of several logics that have been studied extensively in finite model theory. Moreover, we explore connections with descriptive complexity theory and examine the status of certain classical results of model theory in the context of this new framework.

§1. Introduction and summary of results. Finite model theory can be succinctly described as the study of logics on classes of finite structures. Since first-order logic has limited expressive power on classes of finite structures, there has been an in depth investigation of several extensions of first-order logic, each of which embodies a different design principle for enhancing expressive power. More prominent among these extensions are fragments of second-order logic, logics with fixpoint operators, the infinitary logics with a finite number of variables, and, more recently, logics with generalized quantifiers.

There are three main research areas in finite model theory. The first is concerned with the expressive power of logics. More specifically, first-order logic and its extensions are compared in terms of their expressiveness on the class of all finite structures. Therefore, two logics are considered to be *equivalent* if they define the same queries on the class of all finite structures, which means that to separate two logics one has to show that there is some

Received July 1, 1996; revised October 11, 1996.

Most of the results contained in this paper were presented by Ph. Kolaitis in an invited talk during the Annual ASL Meeting at Madison, Wisconsin, in March 1996.

Research of L. Hella was partially supported by a grant from the University of Helsinki.

Research of Ph. Kolaitis was partially supported by NSF Grants No. CCR-9307758 and INT-9024681.

Research of K. Luosto was supported by the Emil Aaltonen Foundation and the Academy of Finland.

© 1996, Association for Symbolic Logic
1079-8986/96/0204-0004/$3.20

class of finite structures on which the two logics have different expressive power. The research in this area has benefited from frequent interaction with database theory and has found applications there (see [1]). Descriptive complexity theory is the second research area in finite model theory; it focuses on the connections between definability in extensions of first-order logic and computational complexity. The prototypical result is Fagin's Theorem [8], which asserts that on the class of all finite structures existential second-order logic Σ^1_1 coincides with NP. Although to this date it is not known if polynomial-time computability can be characterized in terms of definability in some logic on the class of all finite structures (see [14]), it turned out that such fundamental complexity classes as PTIME, NLOGPSACE, and PSPACE are captured by logics with fixpoint operators on the class of all *ordered* finite structures [18, 19, 32, 2]. The third area of finite model theory studies the connections between logic and random structures. It focuses on 0–1 laws for first-order logic and its extensions with respect to various asymptotic probability measures. The origins of this area go back to Glebskii et al. [11] and Fagin [10], who showed independently that first-order logic has a 0–1 law with respect to the uniform measure.

Our goal in this paper is to introduce a new framework for comparing logics in finite model theory. For this, we bring into center stage the concept of *almost everywhere equivalence of logics*, that is to say, two logics having the same expressive power on a class of measure 1. More precisely, if $\mathcal{L}$, $\mathcal{L}'$ are two logics and μ is an asymptotic measure, then $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu)$ means that there is a class $\mathcal{C}$ of structures with $\mu(C) = 1$ and such that $\mathcal{L}$ and $\mathcal{L}'$ define the same queries on $\mathcal{C}$. Although the concept of almost everywhere equivalence seems to underlie much of the work on logic and random structures, so far it has not been studied in its own right. We feel that this concept not only gives rise to a new framework for comparing and classifying logics on finite structures, but also has the potential to enhance the synergy between the three research areas of finite model theory. In particular, we believe that it can create bridges between the areas of expressive power of logics and descriptive complexity theory on the one side and the area of logic and random structures on the other. So far, there has been rather limited interaction between the first two areas and the third one, perhaps because the former have concentrated on fixed classes of finite structures (such as the class of all finite structures and the class of all *ordered* finite structures), while the latter is concerned with the behavior of logics on arbitrary classes of measure 1.

To substantiate the above claims, we embark on a systematic investigation of almost everywhere equivalence of logics and obtain a number of technical results. Note that $\equiv_{\text{a.e.}}$ is an equivalence relation between logics. As a matter of fact, this equivalence relation arises from the preorder relation

$\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}'(\mu)$, which holds when there is a class $\mathcal{C}$ with $\mu(\mathcal{C}) = 1$ and such that $\mathcal{L}'$ is at least as expressive as $\mathcal{L}$ on $\mathcal{C}$. Thus, given two logics, the question is whether they are almost everywhere equivalent and, if not, whether and how they compare in the $\leq_{\text{a.e.}}$ preorder. It should be pointed out that $\mathcal{L} \not\equiv_{\text{a.e.}} \mathcal{L}'(\mu)$ constitutes a strong separation between $\mathcal{L}$ and $\mathcal{L}'$, since it means that $\mathcal{L}$ and $\mathcal{L}'$ have different expressive power on *every* class $\mathcal{C}$ of measure 1.

Our investigation concentrates mainly on almost everywhere equivalence with respect to the uniform measure, which is the most studied one in the area of logic and random structures. Since we are interested in logics that are at least as expressive as first-order logic FO, the $\equiv_{\text{a.e.}}$-equivalence class of FO constitutes the *minimum* member of the preorder $\leq_{\text{a.e.}}$. We establish that several logics of interest in finite model theory are members of this minimum $\equiv_{\text{a.e.}}$-equivalence class. Indeed, we show $\mathcal{L}^{\omega}_{\infty\omega} \equiv_{\text{a.e.}} \text{FO}$, where $\mathcal{L}^{\omega}_{\infty\omega}$ is the infinitary logic that allows for arbitrary disjunctions and conjunctions, but has only a finite number of distinct variables. It follows that FO is also almost everywhere equivalent to logics with fixpoint operators, such as inflationary fixpoint logic and partial fixpoint logic, since the expressive power of these logics is between FO and $\mathcal{L}^{\omega}_{\infty\omega}$.

At the other end of the scale, $\leq_{\text{a.e.}}$ has a *maximum* member, namely, the $\equiv_{\text{a.e.}}$-equivalence class of the infinitary logic $\mathcal{L}_{\infty\omega}$, which can express *every* query on finite structures by allowing for formulae with arbitrary connectives and infinitely many distinct variables. Our main result about the maximum $\equiv_{\text{a.e.}}$-equivalence class is that $\mathcal{L}_{\infty\omega}$ is almost everywhere equivalent to extensions of $\mathcal{L}^{\omega}_{\infty\omega}$ with generalized quantifiers that express *counting* properties, such as the *Even* quantifier D_2 that counts cardinalities modulo 2 and *Rescher's* quantifier R that compares the cardinalities of two sets. The proofs of these results involve a rather unexpected connection between the analysis of the maximum $\equiv_{\text{a.e.}}$-equivalence and the average complexity of algorithms for the *Graph Isomorphism Problem*. Karp [21] and Babai, Erdős, Selkow [3] discovered polynomial-time algorithms that produce canonical labelings on almost all graphs and then used them to obtain algorithms for the *Graph Isomorphism Problem* of polynomial-time average complexity. We show that variants of these canonical labeling algorithms can be defined in first-order logic with the quantifier R and in inflationary fixpoint logic with the quantifier D_2. Therefore these logics can define a linear order on a class of graphs of measure 1. This makes it possible to show that

$$\mathcal{L}^{\omega}_{\infty\omega}(R) \equiv_{\text{a.e.}} \mathcal{L}^{\omega}_{\infty\omega}(D_2) \equiv_{\text{a.e.}} \mathcal{L}_{\infty\omega}.$$

It should be pointed out that on the class of all finite structures $\mathcal{L}_{\infty\omega}$ has strictly higher expressive power than $\mathcal{L}^{\omega}_{\infty\omega}(R)$ and $\mathcal{L}^{\omega}_{\infty\omega}(D_2)$. The analysis of the maximum $\equiv_{\text{a.e.}}$-equivalence class has a number of additional applications to descriptive complexity. In particular, PTIME turns out to be almost everywhere equivalent to the extension of inflationary fixpoint logic with

the quantifier R or the quantifier D_2. This result should be contrasted with the aforementioned open problem of whether there is a logic that captures PTIME on the class of all graphs.

In addition to analyzing the minimum and maximum members of $\leq_{\text{a.e.}}$, we investigate its structural properties as a preorder and show that it has both infinite ascending chains and arbitrarily long finite antichains. Finally, we examine the status of some classical theorems of the model theory of first-order logic in the new framework. As is well known (see [13]), Craig's interpolation theorem, Beth's definability theorem, and the preservation under extensions theorem do *not* hold on the class of all finite structures. We show that this failure of Craig's interpolation theorem and Beth's definability theorem is much deeper, in the sense that they fail on every class $\mathcal{C}$ of finite structures such that $\mu(\mathcal{C}) = 1$. In contrast, we also show there is a class of finite structures of asymptotic measure 1 on which the preservation under extensions theorem holds.

This is an extended abstract that contains only hints or brief sketches of most proofs. We have included, however, extensive background material in Section 2, hoping to make the results presented here accessible to a broad readership.

§2. Background. A *vocabulary* σ is a finite set $\{R_1, \ldots, R_m\}$ of relation symbols of fixed arities. A *σ-structure* $\mathbf{A} = (A, R_1^{\mathbf{A}}, \ldots, R_m^{\mathbf{A}})$ consists of a set A, called the *universe* of $\mathbf{A}$, and relations $R_i^{\mathbf{A}}$ on A whose arities match those of R_i, $1 \leq i \leq m$. In what follows, we shall assume that all structures are finite with universes of the form $\{0, \ldots, n-1\}$ for some positive integer n. For this reason, the term *σ-structure* stands for a finite σ-structure, the term *class $\mathcal{C}$ of σ-structures* means that $\mathcal{C}$ is a class of finite σ-structures that is closed under isomorphisms, and the notation $\text{Str}(\sigma)$ denotes the class of all finite σ-structures.

A *k-ary query Q on* a class $\mathcal{C}$ of σ-structures is a mapping defined on $\mathcal{C}$ and such that if $\mathbf{A}$ is a member of $\mathcal{C}$, then $Q(\mathbf{A})$ is a k-ary relation on A preserved under isomorphisms. A *Boolean query* on a class $\mathcal{C}$ of σ-structures is a mapping Q from $\mathcal{C}$ to $\{0, 1\}$ such that if $\mathbf{A}$ and $\mathbf{B}$ are isomorphic members of $\mathcal{C}$, then $Q(\mathbf{A}) = Q(\mathbf{B})$. Note that a Boolean query on $\mathcal{C}$ can be identified with a subclass of $\mathcal{C}$ closed under isomorphisms. A k-ary *query on σ* is a k-ary query on the class $\text{Str}(\sigma)$ of all σ-structures. The concept of a *Boolean query on σ* is defined in a similar way.

As a representative example of a query on the class $\mathcal{G}$ of all finite graphs, consider the *Transitive Closure* binary query *TC*, which arises by asking: "given a graph $\mathbf{G} = (V, E)$ and two nodes a and b, is there a path from a to b?" More formally, if $\mathbf{G} = (V, E)$ is a graph, then

$$TC(\mathbf{G}) = \{\, (a, b) : \text{there is a path from } a \text{ to } b \,\}.$$

Similarly, the question "given a graph $\mathbf{G} = (V, E)$, is it 3-colorable?" gives rise to the 3-*Colorability* Boolean query Q, where

$$Q(\mathbf{G}) = \begin{cases} 1 & \text{if } \mathbf{G} \text{ is 3-colorable} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\mathcal{L}$ be a logic and Q a query on a class $\mathcal{C}$ of σ-structures. We say that Q is $\mathcal{L}$-*definable* if there is a formula $\psi(x_1, \ldots, x_k)$ of $\mathcal{L}$ such that for every $\mathbf{A} \in \mathcal{C}$

$$Q(\mathbf{A}) = \{ (a_1, \ldots, a_k) : \mathbf{A}, a_1, \ldots, a_k \models \psi(x_1, \ldots, x_k) \}.$$

For instance, neither the *Transitive Closure* query nor the 3-*Colorability* query is first-order definable; both these queries, however, are Σ^1_1-definable.

$\mathcal{L}[\mathcal{C}]$ will denote the class of all $\mathcal{L}$-definable queries on $\mathcal{C}$. Thus, $\mathcal{L}[\mathcal{C}] \subseteq \mathcal{L}'[\mathcal{C}]$ means that every $\mathcal{L}$-definable query on $\mathcal{C}$ is also $\mathcal{L}'$-definable. Finally, $\mathcal{L} \subseteq \mathcal{L}'$ denotes that $\mathcal{L}[\mathrm{Str}(\sigma)] \subseteq \mathcal{L}'[\mathrm{Str}(\sigma)]$ for every vocabulary σ.

2.1. Logics with fixpoint operators. Let $\mathcal{L}$ be a logic and let $\phi(x_1, \ldots, x_n, S)$ be a formula of $\mathcal{L}$ over a vocabulary $\sigma \cup \{S\}$, where S is an n-ary relation symbol that is not in σ. On every σ-structure $\mathbf{A}$, the formula ϕ gives rise to an operator Φ from n-ary relations R on the universe A of $\mathbf{A}$ to n-ary relations on A, where

$$\Phi(R) = \{ (a_1, \ldots, a_n) : \mathbf{A}, a_1, \ldots, a_n, R \models \phi(x_1, \ldots, x_n, S) \}.$$

Such an operator Φ can be iterated by defining the sequence of *stages* Φ^m, $m \geq 0$, where

$$\Phi^0 = \emptyset, \quad \Phi^{m+1} = \Phi(\Phi^m).$$

On a given σ-structure, either there is a positive integer m such that $\Phi^m = \Phi^{m+1}$ or the sequence Φ^m, $m \geq 0$, of stages cycles without ever yielding a fixpoint of Φ. We define the *partial fixpoint* of ϕ to be a stage Φ^m such that $\Phi^m = \Phi^{m+1}$, if such a stage exists, or the empty set $\emptyset$, otherwise.

An operator Φ may be also iterated in an *inflationary* manner, so that the sequence of stages is forced to increase by ensuring that each stage contains the previous stage. More formally, the *inflationary stages* Φ^m_*, $m \geq 0$, of Φ are defined by the induction:

$$\Phi^0_* = \emptyset, \qquad \Phi^{m+1}_* = \Phi^m_* \cup \Phi(\Phi^m_*).$$

The *inflationary fixpoint* of ϕ is the first inflationary stage $\Phi^{m_0}_*$ such that $\Phi^{m_0}_* = \Phi^m_*$ for all $m \geq m_0$.

Partial Fixpoint Logic PFP and *Inflationary Fixpoint Logic* IFP are the extensions of first-order logic that are closed under first-order operations and contain respectively the partial fixpoints and the inflationary fixpoints of all first-order formulae. Gurevich and Shelah [15] established that in terms of expressive power inflationary fixpoint logic IFP coincides with

Least Fixpoint Logic LFP, which augments first-order logic with the least fixpoints of all positive first-order formulae.

It is clear that IFP $\subseteq$ PFP. Moreover, the relationships of IFP and PFP to complexity classes are well understood. First, it is easy to verify that IFP $\subseteq$ PTIME and PFP $\subseteq$ PSPACE. Actually, on the class $\mathrm{Str}(\sigma)$ of all σ-structures these containments are proper, since, as shown in [5], *counting* queries, such as "is there an even number of elements?", are not expressible in PFP. In contrast, on classes of linearly ordered structures IFP captures PTIME [18, 32] and PFP captures PSPACE [32, 2]. More formally, if σ contains a binary relation symbol $<$ and $\mathcal{C}$ is a class of σ-structures such that $<^{\mathbf{A}}$ is a linear order on A for every $\mathbf{A}$ in $\mathcal{C}$, then

$$\mathrm{IFP}[\mathcal{C}] = \mathrm{PTIME}[\mathcal{C}] \quad \text{and} \quad \mathrm{PFP}[\mathcal{C}] = \mathrm{PSPACE}[\mathcal{C}].$$

These two theorems are primary instances of results in descriptive complexity theory, where computational complexity concepts are characterized in terms of definability in logics on ordered finite structures.

2.2. Infinitary logics with a fixed number of variables. The infinitary logic $\mathcal{L}_{\infty\omega}$ is the extension of first-order logic that allows for infinite disjunctions and conjunctions in the syntax. It turns out that $\mathcal{L}_{\infty\omega}$ is too powerful to make any meaningful distinctions in finite model theory, since $\mathcal{L}^{\omega}_{\infty\omega}$ can define *every* query on the class of all finite structures. In general, formulae of $\mathcal{L}_{\infty\omega}$ may have an infinite number of distinct variables. By viewing the number of distinct variables as a resource, we obtain fragments of $\mathcal{L}_{\infty\omega}$ in which the total number of variables is required to be finite.

If k is a positive integer, then the *infinitary logic* $\mathcal{L}^{k}_{\infty\omega}$ *with* k *variables* consists of all formulae of $\mathcal{L}_{\infty\omega}$ with at most k distinct variables. The infinitary logic $\mathcal{L}^{\omega}_{\infty\omega}$ consists of all formulae of $\mathcal{L}_{\infty\omega}$ with a finite number of distinct variables, that is to say, $\mathcal{L}^{\omega}_{\infty\omega} = \bigcup_{k=1}^{\infty} \mathcal{L}^{k}_{\infty\omega}$. It is well known that $\mathcal{L}^{\omega}_{\infty\omega}$ subsumes both IFP and PFP on finite structures (see [25]), that is, $\mathrm{IFP} \subseteq \mathrm{PFP} \subseteq \mathcal{L}^{\omega}_{\infty\omega}$.

2.3. Logics with generalized quantifiers. Generalized quantifiers were introduced by Mostowski [30] who studied first order logic with cardinality quantifiers. Lindström [26] extended the concept so that arbitrary queries can be made definable by adding a generalized quantifier. Lindström's [26] definition adapts easily to finite model theory.

A *generalized quantifier of vocabulary* $\tau = \{T_1, \ldots, T_s\}$ is a Boolean query Q on τ, that is to say, it is a collection of finite structures over τ closed under isomorphisms. The *type* of Q is the sequence $(t_1, \ldots, t_s)$, where $\mathrm{arity}(T_i) = t_i$, $1 \leq i \leq s$. Such a quantifier Q can be applied to any sequence of formulae $\varphi_i(\mathbf{x}_i, \mathbf{y})$, $1 \leq i \leq s$, in which $\mathrm{length}(\mathbf{x}_i) = \mathrm{arity}(T_i)$ for every $i \leq s$. The resulting formula $(Q\,\mathbf{x}_1, \ldots, \mathbf{x}_s)(\varphi_1, \ldots, \varphi_s)$ has the

following semantics:

$$\mathbf{B}, \mathbf{b} \models (Q\,\mathbf{x}_1, \ldots, \mathbf{x}_s)(\varphi_1, \ldots, \varphi_s) \quad \text{if and only if} \quad \mathcal{F}(\mathbf{B}, \mathbf{b}) \in Q,$$

where $\mathcal{F}(\mathbf{B}, \mathbf{b})$ is the τ-structure obtained by interpreting the formulae $\varphi_1, \ldots, \varphi_s$ on $(\mathbf{B}, \mathbf{b})$, i.e.,

$$\mathcal{F}(\mathbf{B}, \mathbf{b}) = (B, \{\, \mathbf{c}_1 : \mathbf{B}, \mathbf{c}_1, \mathbf{b} \models \varphi_1(\mathbf{x}_1, \mathbf{y}) \,\}, \ldots, \{\, \mathbf{c}_s : \mathbf{B}, \mathbf{c}_s, \mathbf{b} \models \varphi_s(\mathbf{x}_s, \mathbf{y}) \,\}).$$

The following two generalized quantifiers will turn out to be of particular use here:

• *Even Quantifier* $D_2 = \{\, (B, X) : (X \subseteq B) \wedge (|X| \text{ is even}) \,\}$. Thus,

$$\mathbf{B} \models (D_2\, x_1)\varphi(x_1) \quad \text{if and only if} \quad |\{\, c_1 : \mathbf{B}, c_1 \models \varphi(x_1) \,\}| \text{ is even.}$$

• *Rescher's Quantifier* $R = \{\, (B, S, T) : (S \subseteq B) \wedge (T \subseteq B) \wedge (|S| < |T|) \,\}$. Thus,

$$\mathbf{B} \models (R\, x_1, x_2)(\varphi(x_1), \psi(x_2)) \quad \text{if and only if} \quad |\{\, c_1 : \mathbf{B}, c_1 \models \varphi(x_1) \,\}| < |\{\, c_2 : \mathbf{B}, c_2 \models \psi(x_2) \,\}|.$$

Let $\mathbf{Q}$ be a set of generalized quantifiers. The logics that we encountered earlier can now be augmented with quantifiers from $\mathbf{Q}$ as follows.

• $\mathrm{FO}(\mathbf{Q})$ and $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q})$ are the extensions of first-order logic and $\mathcal{L}^{\omega}_{\infty\omega}$ that are closed under the respective operations of these logics and applications of quantifiers Q in $\mathbf{Q}$.

• $\mathrm{IFP}(\mathbf{Q})$ and $\mathrm{PFP}(\mathbf{Q})$ are the extensions of $\mathrm{FO}(\mathbf{Q})$ that contain respectively the inflationary fixpoints and the partial fixpoints of all $\mathrm{FO}(\mathbf{Q})$-formulae and are closed under first-order operations and applications of quantifiers Q in $\mathbf{Q}$.

In terms of expressive power, these logics are related in the following way:

$$\mathrm{FO}(\mathbf{Q}) \subseteq \mathrm{IFP}(\mathbf{Q}) \subseteq \mathrm{PFP}(\mathbf{Q}) \subseteq \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q}).$$

2.4. Asymptotic probabilities and 0–1 laws. For every $n \geq 1$, let $\mathrm{Str}_n(\sigma)$ be the set of all σ-structures with universe $\{0, \ldots, n-1\}$. An *asymptotic measure* on the vocabulary σ is a sequence $\mu = (\mu_n)_{n \geq 1}$ of probability measures μ_n on $\mathrm{Str}_n(\sigma)$.

If Q is a Boolean query on σ, then the *asymptotic probability of Q with respect to μ* is

$$\mu(Q) = \lim_{n \to \infty} \mu_n(Q \cap \mathrm{Str}_n(\sigma)),$$

provided this limit exists.

A logic $\mathcal{L}$ has a *convergence law with respect to μ* if $\mu(Q)$ exists for every $\mathcal{L}$-definable query Q on σ. A logic $\mathcal{L}$ has a 0–1 *law with respect to μ* if $\mu(Q) = 0$ or $\mu(Q) = 1$ for every $\mathcal{L}$-definable query Q on σ.

The study of 0–1 laws and convergence laws for various logics has been one of the main research themes in finite model theory. Rather complete answers have been obtained for the asymptotic probabilities of several logics with respect to the *uniform* measure $\mu = (\mu_n)_{n\geq 1}$, which is the measure that assigns the same probability $\mu_n(\{\mathbf{A}\}) = 1/|\operatorname{Str}_n(\sigma)|$ to each structure $\mathbf{A} \in \operatorname{Str}_n(\sigma)$. Glebskii et al. [11] and Fagin [10] established that first-order logic has the 0–1 law with respect to the uniform measure. The 0–1 law holds for $\mathcal{L}^{\omega}_{\infty\omega}$ with respect to the uniform measure [25], but even the convergence law fails for monadic second-order logic [22]. The 0–1 law for $\mathcal{L}^{\omega}_{\infty\omega}$ (and, hence, for first-order logic as well) with respect to the uniform measure can be proved using the *extension axioms* for the vocabulary σ. More precisely, for every $k \geq 1$ the *extension axiom* θ_k is a first-order sentence with k variables asserting that every σ-structure with fewer than k elements can be extended to every possible σ-structure with k elements. Let $\mathcal{C}_k$ be the class of all σ-structures that satisfy θ_k. Fagin [10] showed that $\mu(\mathcal{C}_k) = 1$ for every $k \geq 1$. Moreover, one can show that $\mathcal{L}^{k}_{\infty\omega}$ collapses to the quantifier-free fragment of $\mathcal{L}^{k}_{\omega\omega}$ on $\mathcal{C}_k$, where $\mathcal{L}^{k}_{\omega\omega}$ is the collection of all first-order formulae with at most k variables (see [25]). The 0–1 law for $\mathcal{L}^{\omega}_{\infty\omega}$ with respect to the uniform measure follows immediately from these two facts.

Starting with the seminal paper by Erdős and Rényi [7], combinatorialists investigated in depth the asymptotic probabilities of graph properties with respect to non-uniform measures. These measures arise from *edge probability functions* $p(n)$ that give the probability of an edge between any two distinct nodes of an undirected graph with n nodes. Spencer and Shelah [31] investigated the existence of convergence laws and 0–1 laws for first-order logic with respect to edge probability functions of the form $p(n) = n^{-\alpha}$, where α is a positive real number. This line of research was extended to $\mathcal{L}^{\omega}_{\infty\omega}$ in [27, 29, 28].

§3. Almost everywhere equivalence of logics. The comparative study of the expressive power of logics has occupied a prominent place in finite model theory. So far, much of the research in this area has focused on the comparison of two logics on the class of all finite structures or on the class of all ordered finite structures, where close connections with computational complexity have been established. In what follows, we initiate a different approach to the comparison of logics in finite model theory by investigating how two given logics compare *almost everywhere* with respect to some asymptotic measure.

DEFINITION 3.1. Let $\mathcal{L}$ and $\mathcal{L}'$ be two logics and μ an asymptotic measure on a vocabulary σ.

• We say that $\mathcal{L}'$ is *almost everywhere at least as strong as* $\mathcal{L}$ *with respect to* μ, and write $\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}'(\mu)$, if there is a class $\mathcal{C}$ of σ-structures such that

$\mu(\mathcal{C}) = 1$ and $\mathcal{L}[\mathcal{C}] \subseteq \mathcal{L}'[\mathcal{C}]$, which means that for every formula $\phi(\mathbf{x})$ of $\mathcal{L}$ there is a formula $\psi(\mathbf{x})$ of $\mathcal{L}'$ such that for every structure $\mathbf{B}$ in $\mathcal{C}$

$$\mathbf{B} \models (\forall \mathbf{x})(\phi(\mathbf{x}) \leftrightarrow \psi(\mathbf{x})).$$

• We say that $\mathcal{L}$ is *almost everywhere equivalent to* $\mathcal{L}'$ *with respect to* μ, and write $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu)$, if $\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}'(\mu)$ and $\mathcal{L}' \leq_{\text{a.e.}} \mathcal{L}(\mu)$.
• $\mathcal{L} <_{\text{a.e.}} \mathcal{L}'(\mu)$ means that $\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}'(\mu)$ and $\mathcal{L} \not\equiv_{\text{a.e.}} \mathcal{L}'(\mu)$.

It is clear that $\leq_{\text{a.e.}}$ is a preorder relation and that $\equiv_{\text{a.e.}}$ is an equivalence relation on logics. Note also that $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu)$ simply means that $\mathcal{L}$ and $\mathcal{L}'$ have the same expressive power on a class of measure 1. Next, we consider a seemingly weaker variant of this notion.

DEFINITION 3.2. Let $\mathcal{L}, \mathcal{L}'$ be two logics and let μ be an asymptotic measure on a vocabulary σ.
• $\mathcal{L} \leq_{\text{w.a.e.}} \mathcal{L}'(\mu)$ means that for every formula $\phi(\mathbf{x})$ of $\mathcal{L}$ there is a formula $\psi(\mathbf{x})$ of $\mathcal{L}'$ and a class $\mathcal{C}$ of σ-structures (that depends on $\phi(\mathbf{x})$) such that $\mu(\mathcal{C}) = 1$ and for every structure $\mathbf{B}$ in $\mathcal{C}$

$$\mathbf{B} \models (\forall \mathbf{x})(\phi(\mathbf{x}) \leftrightarrow \psi(\mathbf{x})).$$

• $\mathcal{L} \equiv_{\text{w.a.e.}} \mathcal{L}'(\mu)$ means that $\mathcal{L} \leq_{\text{w.a.e.}} \mathcal{L}'(\mu)$ and $\mathcal{L}' \leq_{\text{w.a.e.}} \mathcal{L}(\mu)$. In this case, we say that $\mathcal{L}$ is *weakly almost everywhere equivalent to* $\mathcal{L}'$ with respect to the measure μ.

Clearly, $\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}'(\mu)$ implies $\mathcal{L} \leq_{\text{w.a.e.}} \mathcal{L}'(\mu)$, and $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu)$ implies $\mathcal{L} \equiv_{\text{w.a.e.}} \mathcal{L}'(\mu)$. The following basic result will help us establish that in many cases the reverse implications hold as well.

LEMMA 3.3. *Assume that for every* $i \geq 1$ *we have a class* $\mathcal{C}_i$ *of* σ*-structures such that* $\mu(\mathcal{C}_i) = 1$. *Then there exists a class* $\mathcal{C}$ *of* σ*-structures such that* $\mu(\mathcal{C}) = 1$ *and for every* $i \geq 1$ *the set difference* $\mathcal{C} \setminus \mathcal{C}_i$ *is finite.*

HINT OF PROOF. There is an increasing sequence $(k_i)_{i \geq 1}$ of natural numbers such that for every $i \geq 1$ and every $n \geq k_i$ we have that

$$\mu_n\left(\left(\bigcap_{j=1}^{i} \mathcal{C}_j\right) \cap \mathrm{Str}_n(\sigma)\right) \geq 1 - 2^{-i}.$$

The desired class $\mathcal{C}$ consists of all σ-structures $\mathbf{A}$ such that, for every i, if $\mathrm{card}(\mathbf{A}) \geq k_i$, then $\mathbf{A} \in \mathcal{C}_i$. ⊣

THEOREM 3.4. *Let* $\mathcal{L}$, $\mathcal{L}'$ *be two logics such that* $\mathcal{L}$ *is countable,* FO $\subseteq \mathcal{L}'$, *and* $\mathcal{L}'$ *is closed under Boolean connectives. If* μ *is an asymptotic measure on* σ, *then*

$$\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}'(\mu) \quad \textit{if and only if} \quad \mathcal{L} \leq_{\text{w.a.e.}} \mathcal{L}'(\mu).$$

HINT OF PROOF. If $\mathcal{L} \leq_{\text{w.a.e.}} \mathcal{L}'(\mu)$, then for each formula $\phi(\mathbf{x})$ of $\mathcal{L}$ there exists a class $\mathcal{C}_\phi$ and a formula $\psi_\phi(\mathbf{x})$ of $\mathcal{L}'$ such that $\mu(\mathcal{C}_\phi) = 1$ and $\phi(\mathbf{x})$ is equivalent to $\psi_\phi(\mathbf{x})$ on $\mathcal{C}_\phi$. Since there are only countably many formulae of $\mathcal{L}$, the preceding Lemma 3.3 implies that there is a class $\mathcal{C}$ of σ-structures such that $\mu(\mathcal{C}) = 1$ and $\mathcal{C} \setminus \mathcal{C}_\phi$ is finite for every $\mathcal{L}$-formula $\phi(\mathbf{x})$. Note that $\phi(\mathbf{x})$ and $\psi_\phi(\mathbf{x})$ need not agree on $\mathcal{C}$. On the other hand, since $\mathcal{C} \setminus \mathcal{C}_\phi$ is finite and every finite structure can be described up to isomorphism in first-order logic, we can find a Boolean combination $\chi_\phi(\mathbf{x})$ of $\psi_\phi(\mathbf{x})$ with finitely many first-order formulae such that $\phi(\mathbf{x})$ and $\chi_\phi(\mathbf{x})$ are equivalent on $\mathcal{C}$. ⊣

COROLLARY 3.5. *Let $\mathcal{L}$, $\mathcal{L}'$ be two logics such that both are countable, at least as expressive as first-order logic, and closed under Boolean connectives. If μ is an asymptotic measure on σ, then*

$$\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu) \quad \textit{if and only if} \quad \mathcal{L} \equiv_{\text{w.a.e.}} \mathcal{L}'(\mu).$$

REMARK 3.6. It should be pointed out that both the concept of almost everywhere equivalence and the concept of weak almost everywhere equivalence are about formulae with free variables, not just sentences. Thus, these concepts are about definable queries of all arities, not just Boolean queries. Grumbach [12] considered a concept of weak almost everywhere equivalence for sentences of two given logics; in general, this concept is less powerful than $\equiv_{\text{w.a.e.}}$.

Note also that most probabilistic properties of logics, including the 0–1 law and the convergence law, are properties of the $\equiv_{\text{a.e.}}$-equivalence class of a logic. In particular, a logic $\mathcal{L}$ has the 0–1 law with respect to μ if and only if every logic $\mathcal{L}'$ such that $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu)$ has the 0–1 law with respect to μ.

Given a logic $\mathcal{L}$ and an asymptotic measure μ, it is natural to ask: what other logics $\mathcal{L}'$ are almost everywhere equivalent to $\mathcal{L}$ with respect to μ? Our main goal is to analyze those equivalence classes of $\equiv_{\text{a.e.}}$ that contain logics of interest in finite model theory. An examination of the uniform measure is the first step in this undertaking.

§4. Almost everywhere equivalence with respect to the uniform measure. We will find out that many results concerning almost everywhere equivalence of logics with respect to the uniform measure do not depend on the underlying vocabulary σ, provided σ is *non-unary*, that is to say, it contains a relation symbol of arity at least 2. For certain results the same proof works on all non-unary vocabularies, while for some other results the proof uses the fact that if σ is non-unary, then the random undirected graph with constant edge probability $p(n) = 1/2$ can be interpreted on random σ-structures with respect to the uniform measure. Indeed, let T be a relation symbol of σ having arity $r > 1$ and let m, n be two positive integers such that $r = k + m$.

If $\theta(x, y)$ is the formula

$$(x \neq y) \wedge (T(\underbrace{x, \ldots, x}_{k}, \underbrace{y, \ldots, y}_{m}) \leftrightarrow T(\underbrace{y, \ldots, y}_{k}, \underbrace{x, \ldots, x}_{m}))$$

and $\mathbf{A}$ is a random σ-structure with respect to the uniform measure, then $\theta(x, y)$ defines on $\mathbf{A}$ a random graph with edge probability $1/2$. In view of the above, we introduce the following notation: $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'$ means that $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}'(\mu)$ holds for the uniform asymptotic measure μ on every non-unary vocabulary σ. The notations $\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}'$ and $\mathcal{L} <_{\text{a.e.}} \mathcal{L}'$ have similar meaning.

4.1. The minimum $\equiv_{\text{a.e.}}$-equivalence class. Since we are interested in logics that are at least as expressive as first-order logic, we have that $\text{FO} \leq_{\text{a.e.}} \mathcal{L}$ for every logic $\mathcal{L}$ considered here. Thus, FO is a member of the *minimum* $\equiv_{\text{a.e.}}$-equivalence class. The next result reveals that many familiar logics inhabit this minimum class.

THEOREM 4.1. *$\mathcal{L}^{\omega}_{\infty\omega}$ is almost everywhere equivalent to* FO *with respect to the uniform measure. Hence,*

$$\text{FO} \equiv_{\text{a.e.}} \text{IFP} \equiv_{\text{a.e.}} \text{PFP} \equiv_{\text{a.e.}} \mathcal{L}^{\omega}_{\infty\omega}.$$

HINT OF PROOF. Let σ be a non-unary vocabulary and for every $k \geq 1$ let $\mathcal{C}_k$ be the class of all σ-structures that satisfy the extension axiom θ_k. As mentioned earlier, for every $k \geq 1$ we have that $\mu(\mathcal{C}_k) = 1$ and $\mathcal{L}^{k}_{\infty\omega}[\mathcal{C}_k] \subseteq \mathcal{L}^{k}_{\omega\omega}[\mathcal{C}_k]$, where $\mathcal{L}^{k}_{\omega\omega}$ is the collection of all first-order formulae with at most k variables. The result follows by applying the basic Lemma 3.3 to the sequence $\mathcal{C}_k$, $k \geq 1$. Note that Theorem 3.4 can not be applied here, since $\mathcal{L}^{\omega}_{\infty\omega}$ is not a countable logic. ⊣

In many respects, Theorem 4.1 offers an *a posteriori* explanation as to why the 0–1 law with respect to the uniform measure holds for first-order logic, for logics with fixpoint operators, and for the infinitary logic $\mathcal{L}^{\omega}_{\infty\omega}$ with finitely many variables.

4.2. Almost everywhere universal order and the maximum $\equiv_{\text{a.e.}}$-equivalence class. The infinitary logic $\mathcal{L}_{\infty\omega}$ is as expressive as any logic, since $\mathcal{L}_{\infty\omega}$ can define every query. Thus, $\mathcal{L} \leq_{\text{a.e.}} \mathcal{L}_{\infty\omega}$ for every logic $\mathcal{L}$, which means that $\mathcal{L}_{\infty\omega}$ is a member of the *maximum* $\equiv_{\text{a.e.}}$-equivalence class. Note also that $\mathcal{L}^{\omega}_{\infty\omega} \not\equiv_{\text{a.e.}} \mathcal{L}_{\infty\omega}$, since $\mathcal{L}^{\omega}_{\infty\omega}$ has the 0–1 law with respect to the uniform measure, while $\mathcal{L}_{\infty\omega}$ does not. Therefore, the minimum $\equiv_{\text{a.e.}}$-equivalence class is different from the maximum one.

What other logics inhabit the maximum $\equiv_{\text{a.e.}}$-equivalence class? It will turn out that the answer to this question is closely connected to the definability of a linear order on classes of measure 1. Before making this connection precise, we need to introduce some additional concepts.

Let σ be a vocabulary and $\mathcal{C}$ a class of σ-structures. A binary query Q on $\mathcal{C}$ is a *universal order on* $\mathcal{C}$ if for every structure $\mathbf{B}$ in $\mathcal{C}$ the value $Q(\mathbf{B})$ of the query Q on $\mathbf{B}$ is a linear order on the universe of $\mathbf{B}$. It is easy to see that a universal order exists on a class $\mathcal{C}$ of σ-structures if and only if every structure in $\mathcal{C}$ is *rigid*, that is, the identity function is the only automorphism of the structure. The definability of a universal order on the class $\mathcal{R}$ of all rigid σ-structures was investigated in [16]. Here, we are interested in the definability of universal orders on classes of asymptotic measure 1.

DEFINITION 4.2. Let Q be a binary query on the class $\mathrm{Str}(\sigma)$ of all σ-structures. We say that Q is an *almost everywhere universal order on* $\mathrm{Str}(\sigma)$ if Q returns a linear order with probability 1 with respect to the uniform measure, i.e.,

$$\lim_{n\to\infty} \mu_n(\{\, \mathbf{B} \in \mathrm{Str}_n(\sigma) : Q(\mathbf{B}) \text{ is a linear order on } \mathbf{B} \,\}) = 1.$$

Every element of a finite linear order is definable by a formula of first-order logic with two variables (see [20]). Using this fact, we can easily obtain the following characterization.

PROPOSITION 4.3. *Let $\mathcal{L}$ be a logic closed under substitution such that $\mathcal{L}^{\omega}_{\infty\omega} \subseteq \mathcal{L}$ and let μ be the uniform measure on a vocabulary σ. The following statements are equivalent*:

- $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}_{\infty\omega}(\mu)$.
- *There is a formula of $\mathcal{L}$ that defines an almost everywhere universal order on* $\mathrm{Str}(\sigma)$.

Thus, the problem of finding extensions $\mathcal{L}$ of $\mathcal{L}^{\omega}_{\infty\omega}$ such that $\mathcal{L} \equiv_{\text{a.e.}} \mathcal{L}_{\infty\omega}$ reduces to the problem of finding extensions of $\mathcal{L}^{\omega}_{\infty\omega}$ that are capable of defining an almost everywhere universal order.

An outstanding algorithmic question is to determine whether the *Graph Isomorphism Problem* is solvable in polynomial time. On the other hand, in the late 1970s researchers found algorithms for the *Graph Isomorphism Problem* of polynomial-time *average* complexity. The key technical step in these algorithms is to show that in effect an almost everywhere universal order on graphs can be computed in polynomial time (combinatorialists use the concept of *canonical labeling*, instead of universal order). The following two results will be of particular interest and use here. In these results, the term *probability of rejection* stands for the probability with respect to the uniform measure on graphs that the algorithm fails to produce a linear order on graphs with n nodes.

THEOREM 4.4 (Babai et al. [3]). *There is an algorithm that has time complexity $O(n^2)$ and computes an almost everywhere universal order on graphs with probability of rejection $O(n^{-1/7})$.*

THEOREM 4.5 (Karp [21]). *There is an algorithm that has time complexity $O(n^2 \log(n))$ and computes an almost everywhere universal order on graphs with probability of rejection $O(n^{3/2}2^{-n/2})$.*

It turns out that variants of the above algorithms for almost universal order can be expressed in extensions of first-order logic FO and inflationary fixpoint logic IFP with generalized quantifiers.

THEOREM 4.6. *Let σ be a non-unary vocabulary.*

- *There is an almost everywhere universal order on* $\mathrm{Str}(\sigma)$ *that is definable in* $\mathrm{FO}(R)$*, where R is Rescher's quantifier.*
- *There is an almost everywhere universal order on* $\mathrm{Str}(\sigma)$ *that is definable in* $\mathrm{IFP}(D_2)$*, where D_2 is the Even quantifier.*

HINT OF PROOF. Since the random graph with edge probability $1/2$ can be interpreted in any random σ-structure, it is enough to show that there exist almost everywhere universal orders on graphs that are definable in $\mathrm{FO}(R)$ and $\mathrm{IFP}(D_2)$.

We say that a node z is a *cut-point* if all nodes u such that $\deg(u) > \deg(z)$ have distinct degrees (call them nodes of *high degree*) and all nodes w such that $\deg(w) \leq \deg(z)$ are uniquely determined by their connections to nodes of high degree. The property "z is a cut-point" is definable in $\mathrm{FO}(R)$. Moreover, there is a formula $\psi(x, y)$ of $\mathrm{FO}(R)$ that defines a linear order on the nodes of every graph that has a cut-point. This formula first asserts that a cut-point z exists and then orders x and y according to their degrees, if at least one of them is a node of high degree, or orders them lexicographically according to their connections to nodes of high degree. The results in Babai et al. [3] imply that a cut-point exists with probability 1. Thus, $\psi(x, y)$ defines an almost everywhere universal order.

The main step in Karp's algorithm [21] is a *splitting procedure* that distinguishes nodes according to the *parity* of their degree with respect to previously defined sets. In the full paper, we show that this splitting procedure is definable in $\mathrm{IFP}(D_2)$. ⊣

From Proposition 4.3 and Theorem 4.6 we derive immediately the following result, which tells that both $\mathcal{L}^{\omega}_{\infty\omega}(R)$ and $\mathcal{L}^{\omega}_{\infty\omega}(D_2)$ are members of the maximum $\equiv_{\text{a.e.}}$-equivalence class.

COROLLARY 4.7. *$\mathcal{L}^{\omega}_{\infty\omega}(R)$ and $\mathcal{L}^{\omega}_{\infty\omega}(D_2)$ are almost everywhere equivalent to $\mathcal{L}_{\infty\omega}$ with respect to the uniform measure, that is,*

$$\mathcal{L}^{\omega}_{\infty\omega}(R) \equiv_{\text{a.e.}} \mathcal{L}^{\omega}_{\infty\omega}(D_2) \equiv_{\text{a.e.}} \mathcal{L}_{\infty\omega}.$$

In contrast, on the class of all finite stuctures $\mathcal{L}_{\infty\omega}$ has strictly higher expressive power than $\mathcal{L}^{\omega}_{\infty\omega}(R)$ and $\mathcal{L}^{\omega}_{\infty\omega}(D_2)$; moreover, $\mathcal{L}^{\omega}_{\infty\omega}(R) \neq \mathcal{L}^{\omega}_{\infty\omega}(D_2)$ (see [24]).

4.3. Almost everywhere descriptive complexity. Immerman [18] and, independently, Vardi [32] showed that on the class of all ordered finite structures least fixpoint logic can express exactly all polynomial-time queries. In contrast, it is an outstanding open problem in finite model theory to determine whether there is a logic that captures PTIME on the class of all finite structures (see Gurevich [14]). Immerman [18] conjectured that if least fixpoint logic is augmented with a mechanism to express *counting* queries (such as *even cardinality*), then it can express all polynomial-time queries. This conjecture, however, was refuted in a decisive manner by Cai, Fürer, and Immerman [4], who showed that $\mathcal{L}^{\omega}_{\infty\omega}$ augmented with all *counting* quantifiers can not express every polynomial-time query. Nevertheless, by combining the preceding Theorem 4.6 with the fact that fixpoint logics capture complexity classes on ordered structures [18, 32, 2], we obtain the following result, which shows that Immerman's conjecture is true in the almost everywhere sense.

COROLLARY 4.8. *Let R be Rescher's quantifier and let D_2 be the Even quantifier. Then*

- $\mathrm{IFP}(R) \equiv_{\mathrm{a.e.}} \mathrm{IFP}(D_2) \equiv_{\mathrm{a.e.}} \mathrm{PTIME}$.
- $\mathrm{PFP}(R) \equiv_{\mathrm{a.e.}} \mathrm{PFP}(D_2) \equiv_{\mathrm{a.e.}} \mathrm{PSPACE}$.

4.4. Almost everywhere hierarchies. So far, we have encountered two distinct $\equiv_{\mathrm{a.e.}}$-equivalence classes, namely, the minimum one and the maximum one. Monadic second order logic MSO gives rise to a third distinct $\equiv_{\mathrm{a.e.}}$-equivalence class. Indeed, $\mathrm{FO} \not\equiv_{\mathrm{a.e.}} \mathrm{MSO}$, because MSO does not have a 0–1 law with respect to the uniform measure [22]. At the same time, $\mathrm{MSO} \not\equiv_{\mathrm{a.e.}} \mathcal{L}_{\infty\omega}$, because MSO is a countable logic, while $\mathcal{L}_{\infty\omega}$ defines uncountably many distinct queries on every class of measure 1. Thus, $\mathrm{FO} <_{\mathrm{a.e.}} \mathrm{MSO} <_{\mathrm{a.e.}} \mathcal{L}_{\infty\omega}$.

In fact, the structure of the preorder $\leq_{\mathrm{a.e.}}$ seems to be rather complex. As evidence, we prove that $\leq_{\mathrm{a.e.}}$ contains infinite ascending chains, as well as arbitrarily long finite antichains.

We first show that the classical time-hierarchy and space-hierarchy theorems for deterministic complexity classes hold also in the almost everywhere sense. Our computational model is the Turing machine with a fixed number of tapes. We use the standard encoding of finite structures by words. Let σ be a vocabulary and fix an order of its symbols. Let $\mathbf{A}$ be a σ-structure and $<$ a linear order that is not necessarily the interpretation of one of the relation symbols of σ. Let $w(\mathbf{A}, <) \in \{0,1\}^*$ be the word beginning with a sequence encoding the cardinality of the universe of $\mathbf{A}$, and continuing with the encodings of the relations of $\mathbf{A}$ in the predescribed order of σ. Note that, according to our convention, the universe of $\mathbf{A}$ is the set $\{0, \ldots, n-1\}$ for

some positive integer n, so that we have a natural order $<$ and the corresponding canonical encoding $w(\mathbf{A}) = w(\mathbf{A}, <)$ in our possession. Note also that the length $m = |w(\mathbf{A})|$ depends only on the cardinality n of the universe of $\mathbf{A}$, say $m = O(p(n))$, for some polynomial $p(n)$.

THEOREM 4.9 (*Time and space hierarchies*).

• *Suppose $s(n) \geq n^2$ and $t(n) \geq n^2$ are two functions such that $t(n)$ is time constructible and $s(n) = o(t(n))$. Then*

$$\mathrm{DTIME}(s) <_{\text{a.e.}} \mathrm{DTIME}(t).$$

• *Suppose $s(n) \geq \log(n)$ and $t(n) \geq \log(n)$ are two space constructible functions such that $s(n) = o(t(n))$. Then*

$$\mathrm{DSPACE}(s) <_{\text{a.e.}} \mathrm{DSPACE}(t).$$

SKETCH OF PROOF. We prove only the time-hierarchy result, the other part being similar. Clearly

$$\mathrm{DTIME}(s(n)) \leq_{\text{a.e.}} \mathrm{DTIME}(t(n)).$$

We have to find a query Q^* that can be computed in time $O(t(n))$, but not in time $O(s(n))$. The classical time-hierarchy result can not be applied directly, because the language it produces need not be the encoding of a query, that is to say, it need not be preserved under isomorphisms. Let $<^{\text{BES}}$ be the almost everywhere universal order on σ-structures produced by the algorithm of Babai et al. [3] (see 4.4). Since $t(n) \geq n^2$, this order can be computed in time $O(t(n))$. Therefore, we can first canonize the input on a class of asymptotic measure 1 by computing $w(\mathbf{A}, <^{\text{BES}}_{\mathbf{A}})$. Next, we use a diagonal argument that mimics the proof of the classical time-hierarchy theorem and obtain a query Q^*. This query Q^* is not almost everywhere equivalent to any query Q computable in time $O(s(n))$, because there is an infinite set S of natural numbers such that Q and Q^* differ on every structure whose cardinality is a member of S. ⊣

Next we strengthen a quantifier hierarchy result in [17]. Recall that every quantifier corresponds to a class of structures for some vocabulary τ. It turns out that τ sets an upper bound for the expressive power of Q. To be more specific, we define the following linear preorder on vocabularies. Let ρ and τ be two vocabularies. Denote by $n_i(\tau)$ the number of i-ary relation symbols in τ. Then $\rho < \tau$, if $n_i(\rho) < n_i(\tau)$ holds for the largest i for which $n_i(\tau) \neq n_i(\rho)$. Note that neither $\rho < \tau$ nor $\tau < \rho$ holds if and only if ρ and τ are the same up to a renaming of symbols.

THEOREM 4.10. *For every non-unary vocabulary τ there is a generalized quantifier Q of vocabulary τ such that for every set $\mathbf{S}$ of generalized quantifiers of vocabularies smaller than τ we have that* $\mathrm{FO}(Q) \not\leq_{\mathrm{a.e.}} \mathrm{PFP}(\mathbf{S})$. *Moreover, Q can be chosen to be in* EXPSPACE.

HINT OF PROOF. The proof is a subtle generalization of the proof of the main result of [17]; it uses diagonalization and counting arguments. We may suppose that $\mathbf{S}$ is the set of all generalized quantifiers of vocabulary smaller than τ. Let $\varphi_i,\ i \geq 0$, be an effective listing of all sentences of $\mathrm{PFP}(\mathbf{S})[\tau]$ up to similarity, that is, sentences φ and φ' are called *similar* if one is obtained from the other by substituting some quantifiers for quantifiers of the same type. We also require that every φ_i occurs infinitely many times in the list. The idea is that φ_is are treated as potential definitions of the quantifier Q. For $n > 0$ and $i \geq 0$, let $f(n, i)$ be the number of sets of the form

$$K_{\varphi_i'} = \{\, \mathbf{A} \in \mathrm{Str}_n(\tau) \mid \mathbf{A} \models \varphi_i' \,\}$$

where φ_i' is similar to φ_i. In [17] it was proved that for each $i > 0$ and all sufficiently large n, we have $f(n, i) < 2^{I(n)}$, where $I(n)$ is the number of isomorpism types of τ-structures of cardinality n. Actually, even more can be shown: for sufficiently large n, we have that $g(n, i) < 2^{I(n)}$, where $g(n, i)$ is the number of $K \subseteq \mathrm{Str}_n(\tau)$ such that for some $K_{\varphi_i'}$ as above $\mu_n(K \,\Delta\, K_{\varphi_i'}) < 1/3$. In conclusion, we can choose an ascending sequence $(n_i)_{i\geq 0}$ of natural numbers and $K_i \subseteq \mathrm{Str}_{n_i}(\tau),\ i \geq 0$, so that any φ_i' similar to φ_i differs from K_i with probability μ_{n_i} at least $1/3$. Let $Q = \bigcup_{i\geq 0} K_i$. Then no $\varphi \in \mathrm{PFP}(\mathbf{S})[\tau]$ defines Q almost everywhere.

As in [17], all steps can be done effectively, including the choice of the sets $K_i,\ i \geq 0$. This particular step can be implemented by testing different sets $K \subseteq \mathrm{Str}_{n_i}(\tau)$ in a canonical order, which requires working space of size $2^{I(n)}$, that is, exponential space. ⊣

Theorem 4.10 implies that $\mathrm{EXPSPACE} \not\equiv_{\mathrm{a.e.}} \mathrm{PFP}(\mathbf{S})$ for every finite set $\mathbf{S}$ of quantifiers. In conjunction with Corollary 4.8, this provides a new intrinsic difference between PSPACE and EXPSPACE. Theorem 4.10 can also be used to establish that the preorder $\leq_{\mathrm{a.e.}}$ contains antichains.

COROLLARY 4.11. *For every positive integer n there exist quantifiers $Q_1, \ldots, Q_n$ such that the logics* $\mathrm{FO}(Q_1), \ldots, \mathrm{FO}(Q_n)$ *are pairwise $\leq_{\mathrm{a.e.}}$-incomparable.*

SKETCH OF PROOF. We consider the case $n = 3$. By Theorem 4.10, there is a quantifier Q_3 of type $(2, 2, 2)$ which is not definable in the extension of first-order logic with quantifiers of type $(2, 2)$; similarly, there is a quantifier S_2 of type $(2, 2)$ which is not definable in the extension of first-order logic with quantifiers of type (2). Since there are uncountably many (in fact, continuum many) pairwise $\equiv_{\mathrm{a.e.}}$-inequivalent queries and since $\mathrm{FO}(Q_3)$ defines only

countably many queries, there is a quantifier S_2' of type $(2,2)$ that is not $\equiv_{\text{a.e.}}$-equivalent to any query definable in $\text{FO}(Q_3)$. By the same token, there is a quantifier Q_1 of type (2) that is not $\equiv_{\text{a.e.}}$-equivalent to any query definable in $\text{FO}(S_2, S_2', Q_3)$. Let Q_2 be a quantifier of type $(2,2,1)$ encoding the pair (S_2, S_2'). Then the logics $\text{FO}(Q_1)$, $\text{FO}(Q_2)$, $\text{FO}(Q_3)$ are pairwise $\leq_{\text{a.e.}}$-incomparable. ⊣

We conclude this section with a diagram that puts together the picture of almost everywhere equivalence of logics with respect to the uniform measure.

$$
\begin{array}{ccccc}
\mathcal{L}^{\omega}_{\infty\omega}(R) & \equiv_{\text{a.e.}} & \mathcal{L}^{\omega}_{\infty\omega}(D_2) & \equiv_{\text{a.e.}} & \mathcal{L}_{\infty\omega} \\
\nearrow & & \uparrow & & \nwarrow \\
\cdots & & \cdots & & \cdots \\
\cdots & & \text{MSO} & & \cdots \\
\cdots & & \cdots & & \cdots \\
\nwarrow & & \uparrow & & \nearrow \\
\text{FO} & \equiv_{\text{a.e.}} & \text{IFP} & \equiv_{\text{a.e.}} & \mathcal{L}^{\omega}_{\infty\omega}
\end{array}
$$

4.5. Almost everywhere model theory. It is well known that practically all main theorems of classical model theory fail on the class of all finite structures, including Craig's interpolation theorem, Beth's definability theorem, and the preservation under extensions theorem (see Gurevich [13]). In this section, we address the question of whether these classical results hold on some class $\mathcal{C}$ of finite structures such that $\mu(\mathcal{C}) = 1$, where μ is the uniform measure. As before, we assume that σ is a non-unary vocabulary.

THEOREM 4.12. *If $\mathcal{C}$ is a class of σ-structures such that $\mu(\mathcal{C}) = 1$, then Beth's definability theorem for first-order logic fails on $\mathcal{C}$.*

SKETCH OF PROOF. Without loss of generality, we may assume that σ contains a binary relation symbol. By Theorem 4.6, there exist a class $\mathcal{C}^*$ of σ-structures and a $\text{FO}(R)$-formula $\psi(x, y)$ such that $\mu(\mathcal{C}^*) = 1$ and $\psi(x, y)$ defines a universal order on $\mathcal{C}^*$. Since Rescher's quantifier R is polynomial-time computable, it is also IFP-definable on every class of ordered finite structures. Let $\prec$ be a binary relation symbol not in the vocabulary σ. It follows that there is an IFP-sentence $\phi(\prec)$ over the vocabulary $\sigma \cup \{\prec\}$ asserting that $\prec$ is a linear order and that $(\forall x)(\forall y)(x \prec y \leftrightarrow \psi(x, y))$. As a result, $\phi(\prec)$ is an IFP-sentence that defines implicitly a universal order on the class $\mathcal{C}^*$. Since, as shown in [23], every IFP-definable query on finite structures is also first-order implicitly definable, it follows that there is a universal order on $\mathcal{C}^*$ that is first-order implicitly definable. Assume now that $\mathcal{C}$ is a class of σ-structures with $\mu(\mathcal{C}) = 1$ and such that Beth's definability theorem for first-order logic holds on $\mathcal{C}$. By combining the preceding facts and hypotheses, we conclude that there is a universal order on $\mathcal{C} \cap \mathcal{C}^*$ that is first-order definable. Consequently, the *even cardinality* query is IFP-definable

on $\mathcal{C} \cap \mathcal{C}^*$ and, since $\mu(\mathcal{C} \cap \mathcal{C}^*) = 1$, its asymptotic probability on $\mathcal{C} \cap \mathcal{C}^*$ must be either 0 or 1. This, however, is a contradiction. ⊣

COROLLARY 4.13. *If $\mathcal{C}$ is a class of σ-structures such that $\mu(\mathcal{C}) = 1$, then Craig's interpolation theorem for first-order logic fails on $\mathcal{C}$.*

The preceding two results reveal that the failure of Craig's interpolation theorem and Beth's definability theorem is inherent to every class of finite structures of asymptotic measure 1. In contrast, the next result shows that the state of affairs is completely different for the preservation under extensions theorem.

THEOREM 4.14. *There is a class $\mathcal{C}$ of σ-structures such that $\mu(\mathcal{C}) = 1$ and every first-order sentence preserved under extensions on $\mathcal{C}$ is equivalent to an existential first-order sentence on $\mathcal{C}$.*

PROOF. As in the proof of Theorem 4.1, for every $k \geq 1$ let $\mathcal{C}_k$ be the class of finite models of the *extension axioms* θ_k. By Lemma 3.3, there is a class $\mathcal{C}$ of σ-structures with $\mu(\mathcal{C}) = 1$ and such that for every $k \geq 1$ the set difference $\mathcal{C} \setminus \mathcal{C}_k$ is finite. We claim that the preservation under extensions theorem holds on $\mathcal{C}$. Let ϕ be a first-order sentence that is preserved under extensions on $\mathcal{C}$. Our task is to show that there is a first-order existential sentence ψ such that ϕ is equivalent to ψ on $\mathcal{C}$. Let m be the total number of distinct variables of ϕ. Since $\mathcal{L}^m_{\infty\omega}$ admits elimination of quantifiers on $\mathcal{C}_m$, either ϕ is true on every member of $\mathcal{C}_m$ or ϕ is false on every member of $\mathcal{C}_m$; in particular, there is an existential sentence ϕ^* such that ϕ is equivalent to ϕ^* on $\mathcal{C}_m$. Let n be the maximum of the cardinalities of structures in $\mathcal{C} \setminus \mathcal{C}_m$, and let $\mathbf{A}_1, \ldots, \mathbf{A}_l$ be an enumeration of all members of $\mathcal{C}$ that have cardinality at most n and satisfy ϕ. For each $j \leq l$, let $\chi_{\mathbf{A}_j}$ be an existential first-order sentence such that a structure $\mathbf{B}$ satisfies $\chi_{\mathbf{A}_j}$ if and only if $\mathbf{B}$ contains a substructure that is isomorphic to $\mathbf{A}_j$. Finally, let ζ_{n+1} be an existential first-order sentence asserting that there exist at least $n+1$ distinct elements. We claim that for every $\mathbf{A}$ in $\mathcal{C}$

$$\mathbf{A} \models \phi \quad \text{if and only if} \quad \mathbf{A} \models \Big(\bigvee_{j=1}^{l} \chi_{\mathbf{A}_j}\Big) \vee (\phi^* \wedge \zeta_{n+1}).$$

Indeed, assume first that $\mathbf{A}$ is a member of $\mathcal{C}$ satisfying ϕ. If $\mathbf{A}$ has at most n elements, then $\mathbf{A}$ is equal to some $\mathbf{A}_j$ and, hence, it satisfies $\chi_{\mathbf{A}_j}$. Otherwise, $\mathbf{A}$ is a member of $\mathcal{C}_m$ and has at least $n+1$ elements, which implies that $\mathbf{A}$ satisfies $\phi^* \wedge \zeta_{n+1}$. Conversely, if $\mathbf{A}$ is a member of $\mathcal{C}$ that satisfies $\chi_{\mathbf{A}_j}$ for some $j \leq l$, then $\mathbf{A}$ satisfies ϕ, since in this case $\mathbf{A}$ contains a substructure isomorphic to $\mathbf{A}_j$ and the sentence ϕ is preserved under extensions on $\mathcal{C}$. If $\mathbf{A}$ is a member of $\mathcal{C}$ that satisfies $\phi^* \wedge \zeta_{n+1}$, then $\mathbf{A}$ must be a member of $\mathcal{C}_m$; thus, it satisfies ϕ as well. ⊣

§5. Almost everywhere equivalence with respect to non-uniform measures. In the full paper, we obtain a number of results about almost everywhere equivalence with respect to non-uniform measures. Here, we only state one result about arbitrary measures and mention some of its applications to random graphs. If μ is a non-uniform asymptotic measure, then $\mathcal{L}^{\omega}_{\infty\omega}$ may not be almost everywhere equivalent to FO with respect to μ. Nevertheless, it is possible to characterize when this happens. Our next theorem complements characterizations of similar flavor for the 0–1 law for $\mathcal{L}^{\omega}_{\infty\omega}$ in [25] and for the convergence law for $\mathcal{L}^{\omega}_{\infty\omega}$ in [29].

THEOREM 5.1. *Let μ be an asymptotic measure on σ. The following statements are equivalent*:

- FO $\equiv_{\text{a.e.}} \mathcal{L}^{\omega}_{\infty\omega}(\mu)$.
- *For every $k \geq 1$, there are finitely many $\mathcal{L}^{k}_{\infty\infty}$-equivalence classes $C_1, \ldots, C_r$ of σ-structures such that $\mu(\bigcup_{i=1}^{r} C_i) = 1$.*
- IFP $\equiv_{\text{a.e.}} \mathcal{L}^{\omega}_{\infty\omega}(\mu)$.

HINT OF PROOF. For every $k \geq 1$, if $\mathcal{C}$ is a class of σ-structures such that $\mathcal{L}^{k}_{\infty\omega}$ collapses to FO on $\mathcal{C}$, then $\mathcal{L}^{k}_{\infty\omega}$ can define only countably many distinct queries on $\mathcal{C}$. It follows that in this case there are finitely many distinct $\mathcal{L}^{k}_{\infty\omega}$-equivalence classes of σ-structures on $\mathcal{C}$. Conversely, if a class $\mathcal{C}$ of σ-structures is the union of finitely many $\mathcal{L}^{k}_{\infty\omega}$-equivalence classes, then $\mathcal{L}^{k}_{\infty\omega}$ collapses to FO on $\mathcal{C}$, since, by results of Dawar, Lindell, and Weinstein [6], every $\mathcal{L}^{k}_{\infty\omega}$-equivalence class is definable by a first-order sentence with k variables. ⊣

The above result is applied in the full paper to study the relationship between FO and $\mathcal{L}^{\omega}_{\infty\omega}$ with respect to asymptotic measures $\mu(\alpha)$ on graphs that arise from edge probability functions of the form $p(n) = n^{-\alpha}$, where α is a positive real number. For example, it turns out that if $\alpha > 1$, then FO $\equiv_{\text{a.e.}} \mathcal{L}^{\omega}_{\infty\omega}(\mu(\alpha))$, while if $\alpha = 1$, then FO $\not\equiv_{\text{a.e.}} \mathcal{L}^{\omega}_{\infty\omega}(\mu(1))$.

§6. Concluding remarks. We brought into center stage the concept of almost everywhere equivalence of logics and initiated a systematic study of its properties. The main feature of this concept is that it gives rise to a novel framework for classifying logics on finite structures and comparing their expressive power. Moreover, it creates bridges between the area of descriptive complexity and the area of asymptotic probabilities of logical properties. Finally, this concept has considerable explanatory power, as it sheds new light to the phenomena of 0–1 laws and convergence laws in finite model theory.

Several problems remain open. A particularly intriguing one is to analyze the $\equiv_{\text{a.e.}}$-equivalence classes of fragments of second-order logic. We conclude by describing a concrete instance of this problem. Fagin [8] established that

existential second-order logic Σ^1_1 captures NP, which implies that separating NP from coNP is equivalent to separating existential second-order logic Σ^1_1 from universal second-order logic Π^1_1. In another seminal paper, Fagin [9] proved that connectivity is definable in monadic Π^1_1, but not in monadic Σ^1_1. Thus, on finite structures monadic Σ^1_1 is different than monadic Π^1_1. We conjecture that a substantially stronger separation holds, namely, monadic Σ^1_1 is *not* almost everywhere equivalent to monadic Π^1_1 with respect to the uniform measure. From a technical point of view, this conjecture appears to be quite challenging to settle. In particular, Fagin's approach can not be used, since almost all graphs are connected.

Acknowledgments. We wish to thank Johan van Benthem who, after listening to a talk based on a preliminary version of this paper, raised the question of whether any of the classical results of model theory, such as Craig's interpolation theorem, hold in the almost everywhere sense.

REFERENCES

[1] S. Abiteboul, R. Hull, and V. Vianu, ***Foundations of databases***, Addison-Wesley, 1995.

[2] S. Abiteboul and V. Vianu, *Fixpoint extensions of first-order logic and Datalog-like languages*, ***Proceedings of the 4th IEEE symposium on logic in computer science***, 1989, pp. 71–79.

[3] L. Babai, P. Erdős, and S. M. Selkow, *Random graph isomorphism*, ***SIAM Journal on Computing***, vol. 9 (1980), pp. 628–635.

[4] J. Cai, M. Fürer, and N. Immerman, *An optimal lower bound on the number of variables for graph identification*, ***Combinatorica***, vol. 12 (1992), no. 4, pp. 389–410.

[5] A. Chandra and D. Harel, *Structure and complexity of relational queries*, ***Journal of Computer and System Sciences***, vol. 25 (1982), pp. 99–128.

[6] A. Dawar, S. Lindell, and S. Weinstein, *Infinitary logic and inductive definability over finite structures*, ***Information and Control***, vol. 119 (1995), no. 2, pp. 160–175.

[7] P. Erdős and A. Rényi, *On the evolution of random graphs*, ***Publications of the Mathematical Institute of the Hungarian Academy of Sciences***, vol. 7 (1960), pp. 17–61.

[8] R. Fagin, *Generalized first-order spectra and polynomial-time recognizable sets*, ***Complexity of computation*** (R.M. Karp, editor), SIAM-AMS Proceedings, vol. 7, 1974, pp. 43–73.

[9] ———, *Monadic generalized spectra*, ***Zeitschrift für Mathematische Logik und Grundlagen der Mathematik***, vol. 21 (1975), pp. 89–96.

[10] ———, *Probabilities on finite models*, ***The Journal of Symbolic Logic***, vol. 41 (1976), pp. 50–58.

[11] Y. V. Glebskii, D. I. Kogan, M. I. Liogonki, and V. A. Talanov, *Range and degree of realizability of formulas in the restricted predicate calculus*, ***Cybernetics***, vol. 5 (1969), pp. 142–154.

[12] S. Grumbach, *A paradox in database theory*, ***3rd ERCIM database research group workshop on updates and constraints handling in advanced database systems, ERCIM-94-W004 CNR***, 1992, pp. 130–143.

[13] Y. Gurevich, *Toward logic tailored for computational complexity*, ***Computation and proof theory*** (M. M. Richter et al., editors), Lecture Notes in Mathematics, no. 1104,

Springer-Verlag, 1984, pp. 175–216.

[14] ———, *Logic and the challenge of computer science*, ***Current trends in theoretical computer science*** (E. Börger, editor), Computer Science Press, 1988, pp. 1–57.

[15] Y. Gurevich and S. Shelah, *Fixed-point extensions of first-order logic*, ***Annals of Pure and Applied Logic***, vol. 32 (1986), pp. 265–280.

[16] L. Hella, Ph. G. Kolaitis, and Luosto. K., *How to define a linear order on finite models*, ***Proceedings of the 9th IEEE symposium on logic in computer science***, 1994, pp. 40–49.

[17] L. Hella, K. Luosto, and J. Väänänen, *The hierarchy theorem for generalized quantifiers*, to appear in the ***The Journal of Symbolic Logic***, 1995.

[18] N. Immerman, *Relational queries computable in polynomial time*, ***Information and Control***, vol. 68 (1986), pp. 86–104.

[19] ———, *Languages that capture complexity classes*, ***SIAM Journal of Computing***, vol. 16 (1987), pp. 760–778.

[20] N. Immerman and D. Kozen, *Definability with bounded number of bound variables*, ***Information and Computation***, vol. 83 (1989), pp. 121–139.

[21] R. Karp, ***Probabilistic analysis of a canonical algorithm for graphs***, Proceedings of Symposia in Pure Mathematics, no. 34, American Mathematical Society, 1979.

[22] M. Kaufmann and S. Shelah, *On random models of finite power and monadic logic*, ***Discrete Mathematics***, vol. 54 (1985), pp. 285–293.

[23] Ph. G. Kolaitis, *Implicit definability on finite structures and unambiguous computations*, ***Proceedings of the 5th IEEE symposium on logic in computer science***, 1990, pp. 168–180.

[24] Ph. G. Kolaitis and J. K. Väänänen, *Generalized quantifiers and pebble games on finite structures*, ***Annals of Pure and Applied Logic***, vol. 74 (1995), no. 1, pp. 23–75.

[25] Ph. G. Kolaitis and M. Y. Vardi, *Infinitary logic and* 0–1 *laws*, ***Information and Computation***, vol. 98 (1992), pp. 258–294, special issue: selections from the fifth annual IEEE symposium on logic in computer science.

[26] P. Lindström, *First order predicate logic with generalized quantifiers*, ***Theoria***, vol. 32 (1966), pp. 186–195.

[27] J. Lynch, *Infinitary logics and very sparse random graphs*, ***Proceedings of the 8th IEEE symposium on logic in computer science***, 1993, pp. 191–198.

[28] J. Lynch and J. Tyszkiewicz, *The infinitary logic of sparse random graphs*, ***Proceedings of the 10th IEEE symposium on logic in computer science***, 1995, pp. 46–53.

[29] M. McArthur, *Convergence and* 0–1 *laws for* $L^k_{\infty\omega}$ *under arbitrary measures*, ***Computer science logic*** (L. Pacholski and J. Tiuryn, editors), Springer-Verlag, 1995, CSL '94, pp. 228–241.

[30] A. Mostowski, *On a generalization of quantifiers*, ***Fundamenda Mathematicae***, vol. 44 (1957), pp. 12–36.

[31] S. Shelah and J. Spencer, *Zero-one laws for sparse random graphs*, ***Journal of the American Mathematical Society***, vol. 1 (1988), pp. 97–115.

[32] M. Y. Vardi, *The complexity of relational query languages*, ***Proceedings of the 14th ACM symposium on theory of computing***, 1982, pp. 137–146.

DEPARTMENT OF MATHEMATICS
P.O. BOX 4
00014 UNIVERSITY OF HELSINKI, FINLAND

E-mail: hella@cc.helsinki.fi

COMPUTER SCIENCE DEPARTMENT
UC SANTA CRUZ
SANTA CRUZ, CALIFORNIA 95064, USA

E-mail: kolaitis@cse.ucsc.edu

DEPARTMENT OF MATHEMATICS
P.O. BOX 4
00014 UNIVERSITY OF HELSINKI, FINLAND

E-mail: kluosto@cc.helsinki.fi