Reflections on Finite Model Theory

Phokion G. Kolaitis*
IBM Almaden Research Center
San Jose, CA 95120, USA
kolaitis@almaden.ibm.com

Abstract

Advances in finite model theory have appeared in LICS proceedings since the very beginning of the LICS Symposium. The goal of this paper is to reflect on finite model theory by highlighting some of its successes, examining obstacles that were encountered, and discussing some open problems that have stubbornly resisted solution.

1 Introduction

During the past thirty years, finite model theory has developed from a collection of sporadic, but influential, early results to a mature research area characterized by technical depth and mathematical sophistication. In this period, finite model theory has been explored not only for its connections to other areas of computer science (most notably, computational complexity and database theory), but also in its own right as a distinct area of logic in computer science. Since the very first LICS Symposium in 1986, LICS has been a natural home for communicating state-of-the-art advances in finite model theory. Moreover, at least five times since its inception in 1995, the Kleene Award for Best Student Paper has been given for work in finite model theory [7, 11, 66, 69, 70].

The invitation to give a talk at LICS 2007 presents an opportunity to reflect on the development and the state of finite model theory today. This paper is not a comprehensive survey of finite model theory. To begin with, space limitations in a conference proceedings make this an impossible task. More importantly, there is no real need for such a survey, given that, by now, there are two books on the subject [26, 62], a monograph on descriptive complexity [48], and a new book with comprehensive overviews of the main topics in finite model theory and its applications [37]. At the same time, this paper is not a "personal perspective" [31] on the development of finite model theory either. Instead, it is an attempt to highlight some of the successes of finite model theory, to examine some of the obstacles that were encountered, and to discuss some open problems that have stubbornly resisted solution.

2 Early Beginnings

In the first half of the 20th Century, finite models were used as a tool in the study of Hilbert's *Entscheidungsproblem*, also known as the *classical decision problem*, which is the satisfiability problem for first-order logic: given a first-order sentence, does it have a model? Indeed, even before this problem was shown to be undecidable by Turing and Church, logicians had identified decidable fragments of first-order logic, such as the Bernays-Schönfinkel Class of all $\exists^*\forall^*$ sentences and the Ackermann Class of all $\exists^*\forall\exists^*$ sentences. The decidability of these two classes was established by proving that the *finite model property* holds for them: if a sentence in these classes has a model, then it has a finite model (see [14, Chapter 6] for a modern exposition of these results). After hard toil over many years, it turned out that these two are the only quantifier prefix classes of first-order logic (with equality) over relational vocabularies for which the satisfiability problem is decidable [14]. A third important class is the Gödel Class of all $\exists^*\forall\forall\exists^*$ sentences. The equality-free fragment of this class has the finite model property, hence it is decidable [35]; in contrast, the full Gödel Class (with equality) is undecidable [36].

Trakhtenbrot's Theorem [75] is generally regarded as the first important result in finite model theory during the second half of the 20th Century. Here, finite models are the object of study, as this result is about *finitely valid* first-order sentences, i.e., first-order sentences true on all finite models.

Theorem 1 *Let σ be a relational vocabulary containing a non-unary relation symbol. The set of all finitely valid first-order sentences over σ is not recursively enumerable.*

Trakhtenbrot's Theorem says that there is *no* effective axiomatization of the set of all finitely valid first-order sentences. It contrasts sharply with Gödel's Completeness Theorem about the set of all valid first-order sentences, and it

*On leave from UC Santa Cruz.

can be construed as an "anti-completeness" theorem for the set of all finitely valid first-order sentences.

The development of finite model theory was also influenced by the quest to resolve certain problems that were articulated in the 1950s. A set S of positive integers is said to be a *spectrum* if there is a first-order sentence ψ such that $S = \{\text{m: } \psi \text{ has a finite model with } m \text{ elements}\}$. Scholz [71] in 1952 and Asser [6] in 1955 posed the following problems about spectra.

Problem 1 The Spectrum Problem.

- (Scholz) *Characterize all spectra.*
- (Asser) *Are spectra closed under complement? In other words, is the complement of a spectrum also a spectrum?*

A problem of a different character was motivated by the preservation-under-substructures theorem of Łoś-Tarski, which asserts that if a first-order sentence ψ is preserved under substructures on all (finite and infinite) models, then there is a universal first-order sentence ψ^* such that ψ is logically equivalent to ψ^*. In 1958, Scott and Suppes [72] asked whether the preservation-under-substructures theorem holds in the finite, and conjectured that it does.

Conjecture 1 (Scott and Suppes) *If a first-order sentence ψ is preserved under substructures on all finite models, then there is a universal first-order sentence ψ^* such that ψ is equivalent to ψ^* on all finite models.*

We will discuss the status of the Spectrum Problem and the Scott-Suppes Conjecture in later sections.

3 Main Themes in Finite Model Theory

The traditional focus of mathematical logic has been the study of logics on the class of all (finite and infinite) structures or on a fixed infinite structure of mathematical significance. The Completeness Theorem and the Compactness Theorem for first-order logic are two key results in the first category. Gödel's Incompleteness Theorem and Tarski's Theorem about elimination of quantifiers on the reals are two key results in the second category, since these are about first-order logic on the structure $\mathcal{N} = (N, +, \times)$ of the integers and on the structure $\mathcal{R} = (R, +, \times)$ of the reals.

In contrast, finite model theory focuses on the study of logics on classes of finite structures. In addition to first-order logic, various other logics have been explored in the context of finite model theory; they include fragments of second-order logic, logics with fixed-point operators, infinitary logics, and logics with generalized quantifiers. These logics have been investigated in numerous classes of finite structures, including the class of all finite graphs, the class of all finite ordered graphs, the class of all planar graphs, the class of graphs of treewidth bounded by some fixed constant, and the class of all finite strings.

After a number of pioneering results obtained in the late 1960s and the 1970s [28, 30, 34, 50], finite model theory was pursued in its own right in the 1980s and beyond. It turned out that new phenomena emerge, when one focuses on classes of finite structures; these phenomena gave finite model theory its own distinct character and set it apart from other areas of mathematical logic. At the same time, finite model theory benefitted from a continuous interaction with certain areas of computer science, especially computational complexity and database theory.

Research in finite model theory has branched into four areas. The first is the study of the connections between computational complexity and uniform definability in logics on finite structures, an area that is known as *descriptive complexity*. The second (and closely related to the first) is the study of the expressive power of logics on finite structures: what can and what cannot be expressed in various logics on classes of finite structures? The third is the study of the connections between logic and asymptotic combinatorics; here the focus is on 0–1 laws and convergence laws for the asymptotic probabilities of sentences of various logics on classes of finite structures. The final area is the study of classical model theory in the finite: do the classical results of model theory (eg., the various preservation theorems, Craig's Interpolation Theorem) hold in the finite?

In what follows, we highlight some of the achievements in these areas, but also comment on some of the obstacles encountered and on certain problems that still remain open. We assume that σ is a non-empty relational vocabulary; we will write $\mathcal{F}$ to denote the class of all finite σ-structures.

3.1 Descriptive Complexity and Expressive Power

Let $\mathcal{C}$ be a class of finite σ-structures and let k be a positive integer. A k-ary *query on* $\mathcal{C}$ is a mapping Q defined on $\mathcal{C}$ and such that if $\mathbf{A}$ is a structure in $\mathcal{C}$, then $Q(\mathbf{A})$ is a k-ary relation on the universe of $\mathbf{A}$ that is invariant under isomorphisms, i.e., if $f : \mathbf{A} \rightarrow \mathbf{B}$ is an isomorphism, then $Q(\mathbf{B}) = f(Q(\mathbf{A}))$. A *Boolean query on* $\mathcal{C}$ is a mapping from $\mathcal{C}$ to $\{0, 1\}$ that is invariant under isomorphisms; if $Q(\mathbf{A}) = 1$, then we say that $\mathbf{A}$ *satisfies* Q, and write $\mathbf{A} \models Q$. Queries formalize and generalize the concept of a decision problem on a class of finite structures, such as CONNECTIVITY and 3-COLORABILITY. From a computational standpoint, we are interested in determining the computational complexity of a given query q. From a logical standpoint, we are interested in determining whether a given query q is (uniformly) *definable in a logic* L of interest, that is to say, whether there is a formula $\varphi(\mathbf{x})$ of L such that, for every $\mathbf{A} \in \mathcal{C}$, we have that $Q(\mathbf{A}) = \{\mathbf{a} : \mathbf{A} \models \varphi(\mathbf{a})\}$.

Descriptive complexity investigates the interplay between computational complexity and logical definability.

First-order logic FO has extremely high expressive power on the structure $\mathbf{N} = (N, +, \times)$ of the integers, since, for instance, every recursive set is FO-definable on $\mathbf{N}$. In contrast, first-order logic has severely limited expressive power on the class $\mathcal{G}$ of all finite graphs, since, as we will see later on, such basic polynomial-time queries as CONNECTIVITY (given a graph, is it connected?) and EVEN CARDINALITY (given a graph, does it have an even number of nodes?) are *not* FO-definable on $\mathcal{G}$. This state of affairs dictates that logics strictly more expressive than first-order logic have to be used in order to express computationally interesting queries on finite structures.

Existential Second-Order Logic and NP Second-order logic SO augments the syntax of first-order logic with second-order quantifiers $\exists R$ and $\forall R$, where R is a n-ary relation symbol, for some $n \geq 1$. Existential second-order Logic ESO is one of the syntactically simplest fragments of SO; it consists of all formulas of the form $\exists R_1 \cdots \exists R_k \varphi$, where $R_1, \ldots, R_k$ are relation symbols of various arities and φ is a first-order formula. For example, the query 2-COLORABILITY is definable by the ESO-formula $\exists R \forall x \forall y (E(x, y) \rightarrow (R(x) \leftrightarrow \neg R(y))$ on the class $\mathcal{G}$ of all finite graphs. Moreover, for each $k \geq 3$, the NP-complete query k-COLORABILITY is ESO-definable on $\mathcal{G}$ by a similar formula. The next result, due to Fagin [28] and known as Fagin's Theorem, tells that it is not an accident that ESO can express NP-complete problems on finite graphs.

Theorem 2 *Let Q be a query on the class $\mathcal{F}$ of all finite σ-structures. Then the following statements are equivalent.*

- *Q is in* NP.
- *Q is* ESO-*definable on $\mathcal{F}$.*

In symbols, $\mathrm{NP} = \mathrm{ESO}$ *on $\mathcal{F}$.*

Theorem 2 is the prototypical result of descriptive complexity. It yields a machine-independent characterization of NP, and reinforces the unity of logic and computation. An outstanding open problems in computational complexity is whether or not NP is closed under complement, that is, whether or not $\mathrm{NP} = \mathrm{coNP}$. Theorem 2 and the NP-completeness of 3-COLORABILITY imply that this problem can be cast as a problem in finite model theory.

Corollary 3 *The following statements are equivalent.*

- NP *is closed under complement, i.e.,* $\mathrm{NP} = \mathrm{coNP}$.
- ESO *is closed under complement on the class $\mathcal{G}$ of all finite graphs.*
- NON-3-COLORABILITY *is* ESO-*definable on $\mathcal{G}$.*

Fagin was motivated by the Spectrum Problem. Note that if a set S of integers is the spectrum of a first-order sentence ψ over a vocabulary $\sigma = \{R_1, \ldots, R_k\}$, then S coincides with the set of all finite models of the ESO-sentence $\exists R_1 \cdots \exists R_k \psi$ over the empty vocabulary. Fagin [28] viewed ESO-formulas over non-empty vocabularies as defining *generalized spectra*, and used them to characterize NP. His proof of Theorem 2 yields also the following characterization of spectra in terms of the complexity class NEXPTIME of problems solvable by a non-deterministic Turing machine in polynomial time; this characterization was obtained independently by Jones and Selman [50].

Theorem 4 *The following statements are equivalent for a set S of positive integers in binary notation.*

- *S is a spectrum.*
- *S is in* NEXPTIME.

Hence, the complement of every spectrum is a spectrum if and only if NEXPTIME *is closed under complement.*

Theorem 4 provides an answer to Scholz's problem to characterize all spectra. It also reveals that Asser's problem is equivalent to a complexity-theoretic problem. Nonetheless, it does not resolve Asser's problem, since determining whether or not NEXPTIME is closed under complement is still an open problem in computational complexity.

Least Fixed-Point Logic and P As time passed, research in descriptive complexity focused on other major complexity classes and, in particular, on the class P of problems solvable by a deterministic Turing machine in polynomial time. To this end, logics with *fixed-point* operators were investigated in depth; such operators augment the expressive power of first-order logic on finite structures by adding a mechanism that embodies recursion.

Let $\varphi(x_1, \ldots, x_k, S)$ be a first-order formula over the vocabulary $\sigma \cup \{S\}$, where S is a k-ary relation symbol not in σ. On every σ-structure $\mathbf{A}$, this formula gives rise to an operator $\Phi : \mathcal{P}(A^k) \rightarrow \mathcal{P}(A^k)$ such that $\Phi(T) = \{(a_1, \ldots, a_k) : \mathbf{A} \models \varphi(a_1, \ldots, a_k, T)\}$, where A is the universe of $\mathbf{A}$ and $\mathcal{P}(A^k)$ is the powerset of A^k. By iterating Φ any finite number of times, we obtain the (finite) *stages* of Φ; formally, they are defined by the induction:

$$
\begin{aligned}
\varphi^1 &= \Phi(\emptyset); \\
\varphi^{n+1} &= \Phi(\varphi^n).
\end{aligned}
$$

Suppose now that the formula $\varphi(x_1, \ldots, x_k, S)$ is *positive in S*, which means that every occurrence of S is within the scope of an even number of negation symbols. This syntactic property of φ implies that the associated operator Φ is *monotone in* S: if $T \subseteq T'$, then $\Phi(T) \subseteq \Phi(T')$. In turn, this implies that the sequence of finite stages is increasing:

$$\varphi^1 \subseteq \varphi^2 \subseteq \ldots \subseteq \varphi^n \subseteq \varphi^{n+1} \subseteq \ldots \quad .$$

It follows that on every structure $\mathbf{A}$, the operator Φ has a least fixed-point φ^∞, that is, there is a smallest relation P such that $\Phi(P) = P$. Moreover, if $\mathbf{A}$ is a finite structure, then there is a positive integer $s \leq |A|^k$, where $|A|$ is the size of the universe of $\mathbf{A}$, such that the least fixed-point φ^∞ of Φ on $\mathbf{A}$ is equal to the finite stage φ^s. Hence, on the class of all finite σ-structures, φ^∞ is definable by the infinitary disjunction $\bigvee_{n\geq 1} \varphi^n$.

Least fixed-point logic LFP is the extension of first-order logic FO obtained by augmenting the syntax of FO with the least fixed-points of positive formulas. As an example, the CONNECTIVITY query is LFP-definable on the class $\mathcal{G}$ of all finite graphs. Indeed, let $\varphi(x, y, S)$ be the positive-in-S first-order formula $E(x, y) \vee (\exists z)(E(x, z) \wedge S(z, y))$. Then, for every $n \geq 1$, we have that the n-th stage φ^n of $\varphi(x, y, S)$ defines the query "there is a path of length at most n from x to y". Consequently, the least fixed-point $\varphi^\infty(x, y)$ of $\varphi(x, y, S)$ defines the TRANSITIVE CLOSURE query. Hence, a graph $\mathbf{G} = (V, E)$ is connected if and only if $\mathbf{G} \models (\forall x)(\forall y)\varphi^\infty(x, y)$. This shows that LFP is strictly more expressive than FO on the class $\mathcal{G}$ of all finite graphs.

LFP is a robust logic, because it is as expressive as *inflationary fixed-point logic* IFP, whose syntax allows for a more relaxed and seemingly more powerful iteration mechanism. Specifically, if $\varphi(\mathbf{x}, S)$ is an arbitrary first-order formula (that is, $\varphi(\mathbf{x}, S)$ need not be positive in S), then the *inflationary stages* of $\varphi(\mathbf{x}, S)$ are defined by the induction:

$$\begin{aligned} \varphi^1 &= \Phi(\emptyset); \\ \varphi^{n+1} &= \varphi^n \cup \Phi(\varphi^n). \end{aligned}$$

Obviously, the sequence of the inflationary stages of $\varphi(\mathbf{x}, S)$ is increasing by definition; it follows that, for every finite structure $\mathbf{A}$, there is a positive integer $m \leq |A|^k$ such that $\varphi^m = \varphi^{m+1}$. This is the *inflationary fixed-point* of $\varphi(\mathbf{x}, S)$ on $\mathbf{A}$; note that if $\varphi(\mathbf{x}, S)$ is positive in S, then the least fixed-point and the inflationary fixed-point coincide. The infinitary disjunction $\bigvee_{n\geq 1} \varphi^n$ defines the inflationary fixed-point of $\varphi(\mathbf{x}, S)$ on all finite structures. *Inflationary fixed-point logic* IFP is the extension of FO obtained by augmenting the syntax of FO with the inflationary fixed-points of arbitrary formulas. Clearly, LFP $\subseteq$ IFP on $\mathcal{F}$. Gurevich and Shelah [44] showed that the reverse inclusion holds.

Theorem 5 *Let Q be a query on the class $\mathcal{F}$ of all finite σ-structures. Then the following statements are equivalent.*

- *Q is* LFP*-definable on $\mathcal{F}$.*
- *Q is* IFP*-definable on $\mathcal{F}$.*

In symbols, LFP = IFP *on $\mathcal{F}$.*

We saw that, on every finite structure, the least fixed-point of a positive formula can be obtained by iterating the formula at most polynomially-many times in the size of the structure. This implies that every LFP-definable query on $\mathcal{F}$ is computable in polynomial time; in symbols, LFP $\subseteq$ P on $\mathcal{F}$. Moreover, LFP can express P-complete problems on finite structures. For example, if $\varphi(x, S)$ is the formula $(A(x) \vee \exists y \exists z(R(x, y, z) \wedge S(y) \wedge S(z)))$, then the least fixed point φ^∞ of $\varphi(x, S)$ expresses the PATH SYSTEMS query, which was shown by Cook [21] to be P-complete under logspace reductions. Nonetheless, LFP is properly contained in P on $\mathcal{F}$, because LFP cannot express simple counting queries on $\mathcal{F}$, such as the EVEN CARDINALITY query [16]. Immerman [47] and Vardi [76] independently showed that this deficiency of LFP can be overcome if a "built-in" linear order on the universe of the structure is available for use in LFP-formulas. More precisely, an *ordered* finite structure is a finite structure of the form $\mathbf{A} = (A, <, R_1, \ldots, R_m)$, where $<$ is a linear order on the universe A of $\mathbf{A}$. We write $\mathcal{O}$ to denote the class of all finite ordered structures over the vocabulary $\{<, R_1, \ldots, R_m\}$; here, for simplicity, we use the same notation for both the relation symbols in the vocabulary and the relations interpreting them on a structure. We can now give the precise statement of the Immerman-Vardi Theorem.

Theorem 6 *Let Q be a query on the class $\mathcal{O}$ of all finite ordered structures. The following statements are equivalent.*

- *Q is in* P.
- *Q is* LFP*-definable on $\mathcal{O}$.*

In symbols, P = LFP *on $\mathcal{O}$.*

The Quest for a Logic for P Research in descriptive complexity has shown that essentially all major computational complexity classes can be characterized in terms of (uniform) definability in various logics on finite structures; a detailed exposition of these results can be found in the monograph [48]. A perusal of these results reveals that they fall into two distinct groups. In the first group, a complexity class is captured by a certain logic on the class of *all* finite structures over some vocabulary; for example, Fagin's characterization of NP is a result of this first type. In the second group, however, a complexity class is captured by a certain logic on the class of all *ordered* finite structures; for example, the Immerman-Vardi characterization of P is a result of this second type. In the second group of results, a linear order is used in the formulas of the logic at hand to show that the computations of the resource-bounded (deterministic or non-deterministic) Turing machines that define the complexity class can be simulated by the formulas of the logic. Note that a linear order is needed in the proof of

Fagin's Theorem as well, but in that case the linear order does not have to be given explicitly; instead, an existential second-order quantifier can be used to guess some linear order on the input. As a matter of fact, this is the first step in showing that every NP problem can be expressed by an existential second-order formula on all finite structures.

Is the explicit presence of a linear order necessary in characterizing certain computational complexity classes using logic? In particular, is there a descriptive-complexity characterization of P on the class of *all* finite structures? These and other related questions were raised quite early, but remain essentially unanswered to this date. In particular, Gurevich [43] made the following bold conjecture.

Conjecture 2 (Gurevich) *There is* no *logic that captures* P *on the class of all finite structures.*

This conjecture, of course, requires that the notion of "a logic that captures polynomial time" be made precise. Gurevich [43] provided a rigorous definition of this notion, which stipulates, among other things, that the logic possess an effective syntax. Prior to Gurevich's work, Chandra and Harel [16] had raised the following closely related problem.

Problem 2 (Chandra and Harel) *Is there an effective enumeration of all polynomial-time computable queries on the class of all finite structures?*

Clearly, if $\mathrm{P} = \mathrm{NP}$, then there is a logic that captures P on the class of all finite structures, since, by Fagin's Theorem, existential second-order logic ESO is such a logic; furthermore, in this case, the effective syntax of ESO yields an effective enumeration of all polynomial-time queries on the class of all finite structures. This means that Gurevich's Conjecture cannot be confirmed (and also that Chandra and Harel's Problem cannot be answered in the negative) without showing at the same time that $\mathrm{P} \neq \mathrm{NP}$.

Numerous subsequent investigations in finite model theory were directly motivated by the quest to make progress towards Conjecture 2 and Problem 2. In particular, these investigations include the systematic study of logics with powerful *generalized quantifiers*, such as *counting quantifiers* (see [15, 45, 54, 62, 67]). As it turned out, none of the logics with generalized quantifiers considered thus far succeeds in capturing polynomial-time computability on the class of all finite structures. On the other hand, Dawar [22] showed that if there is a logic for P, then there is one that is an extension of first-order logic with a uniform sequence of generalized quantifiers expressing a P-complete problem. These and several other concerted efforts notwithstanding, Chandra and Harel's Problem and Gurevich's Conjecture remain outstanding open problems.

As finite model theory continued to develop, there was a shift in emphasis from the class of all finite structures to restricted classes of finite structures of combinatorial or graph-theoretic interest. In some of these restricted classes, it is possible to use deeper properties of the structures to define a linear order in LFP or in certain extensions of LFP; in turn, this yields a logic for P on such classes. For instance, Immerman and Lander [49] and, independently, Lindell [63] showed that $(\mathrm{IFP} + \mathrm{C})$, the extension of IFP with counting quantifiers, captures P on the class of all finite trees. This result was vastly generalized in two different directions by Grohe [40] and by Grohe and Mariño [41].

Theorem 7 *Let* $(\mathrm{IFP} + \mathrm{C})$ *be the extension of inflationary fixed-point logic* IFP *with counting quantifiers.*

- (Grohe [40]) $\mathrm{P} = (\mathrm{IFP} + \mathrm{C})$ *on the class of all planar graphs.*
- (Grohe and Mariño [41]) *Let* k *be a positive integer and let* $\mathcal{T}(k)$ *be the class of all graphs of treewidth at most* k. *Then* $\mathrm{P} = (\mathrm{IFP} + \mathrm{C})$ *on* $\mathcal{T}(k)$.

Treewidth measures how "tree-like" a graph is. Note that trees are planar graphs and have treewidth 1; planar graphs can have arbitrarily large treewidth, as the $k \times k$-grid is planar and has treewidth k. For background material on the concept of treewidth and its many uses in graph theory and computational complexity, we refer the reader to [25, 33].

Partial Fixed-Point Logic and PSPACE So far, our discussion of descriptive complexity has focused on the complexity classes P and NP. We conclude this section by outlining a descriptive-complexity characterization of PSPACE that will be of interest to us later on.

Let $\varphi(\mathbf{x}, S)$ be a first-order formula that is not necessarily positive in S. Observe that the sequence φ^n, $n \geq 1$, of the finite stages of $\varphi(\mathbf{x}, S)$ is still well defined, but this sequence need not be an increasing one, and may not converge to a fixed point of $\varphi(\mathbf{x}, S)$ (in fact, if $\varphi(\mathbf{x}, S)$ is not positive in S, then it may have no fixed-points whatsoever). Abiteboul and Vianu [3] defined the *partial fixed-point* of $\varphi(\mathbf{x}, S)$ as follows. If $\mathbf{A}$ is a finite structure, then

$$\varphi^\infty = \begin{cases} \varphi^m & \text{if } \mathbf{A} \models \varphi^m = \varphi^{m+1} \text{ for some } m \\ \emptyset & \text{otherwise} \end{cases}$$

Note that if $\varphi(\mathbf{x}, S)$ is positive in S, then its partial fixed-point coincides with its least fixed-point.

Abiteboul and Vianu [3] introduced *partial fixed-point logic* PFP as the extension of first-order logic FO obtained by augmenting the syntax of FO with the partial fixed-points of arbitrary formulas. It is easy to see that LFP is contained in PFP on finite structures and that every PFP-definable query is in PSPACE. Thus, the containments $\mathrm{LFP} \subseteq \mathrm{PFP} \subseteq \mathrm{PSPACE}$ hold on $\mathcal{F}$. Note also that PFP can express PSPACE-complete queries on finite structures. For this, let $\psi(x, S)$ be the first-order formula

$(A(x) \vee \exists y \exists z(S(y) \wedge \neg S(z) \wedge R(x,y,z)))$. Grohe [39] showed that the partial fixed-point ψ^∞ of $\psi(x,S)$ defines a PSPACE-complete problem. Nonetheless, PFP suffers from the same deficiencies as LFP when it comes to expressing counting queries. In particular, EVEN CARDINALITY is not PFP-definable on $\mathcal{F}$. The state of affairs, however, is different if a linear order is explicitly available.

Theorem 8 *Let Q be a query on the class $\mathcal{O}$ of all finite ordered structures. The following statements are equivalent.*

- *Q is in* PSPACE*.*
- *Q is* PFP*-definable on $\mathcal{O}$.*

In symbols, $\mathrm{PSPACE} = \mathrm{PFP}$ *on $\mathcal{O}$.*

Chandra and Harel [16] introduced $(\mathrm{FO} + \mathrm{While})$ as an extension of FO with "while looping" as a mechanism for recursion. Theorem 8 follows by combining a result of Vardi [76] that $\mathrm{PSPACE} = (\mathrm{FO} + \mathrm{While})$ on $\mathcal{O}$ with a result of Abiteboul and Vianu [3] that $(\mathrm{FO}+\mathrm{While}) = \mathrm{PFP}$.

3.2 Games and the Expressive Power of Logics

The results in descriptive complexity highlighted in the previous section provide a way to calibrate the expressive power of logics on classes of finite structures by matching it against a complexity class. In most cases, however, these results do not help in determining whether the expressive power of one logic is provably different from that of another logic, since separating two logics in the finite often amounts to separating two complexity classes. Consider, for instance, the question of whether or not existential second-order logic is different from universal second-order logic on the class of finite structures (i.e., whether existential second-order logic is closed under complement in the finite). As we saw in Corollary 3, this question is equivalent to whether or not $\mathrm{NP} = \mathrm{coNP}$, which is an outstanding open problem.

Combinatorial games provide a tool for analyzing and delineating the expressive power of logics on classes of finite structures. The most well-studied such games are the Ehrenfeucht-Fraïssé games for first-order logic FO. They yield a sound and complete method for investigating the expressive power of FO. Moreover, Ehrenfeucht-Fraïssé games are an adaptable and versatile tool, as the expressive power of several logics stronger than FO can be analyzed using suitable variants of these games.

These combinatorial games are played between two players, called the *Spoiler* and the *Duplicator*, on two structures. The players take turns and, depending on the game, choose elements from the structures, or relations on the structures, or place pebbles on elements of the structures. Intuitively, the goal of the Spoiler is to establish a "difference" between the two structures, while the goal of the Duplicator is to maintain "similarity". Such a game G is tailored for a logic L on a class $\mathcal{C}$ of σ-structures if the following statements are equivalent for a Boolean query Q on the class $\mathcal{C}$:

- Q is not definable on $\mathcal{C}$ by a sentence of L.
- There are structures $\mathbf{A}$ and $\mathbf{B}$ in $\mathcal{C}$ such that $\mathbf{A} \models \mathbf{Q}$, $\mathbf{B} \not\models Q$, and the Duplicator wins the game G on $\mathbf{A}$ and $\mathbf{B}$.

If the above holds, then the game G gives rise to a sound and complete methodology for analyzing definability in L. Typically, the game used is not a single game, but, rather, a family of similar games parameterized by some parameter (eg., number of moves or number of pebbles) that corresponds to some resource in the syntax of the logic (eg., quantifier depth or number of variables). For precise definitions of the various games mentioned here, we refer the reader to [26, 37, 48, 62].

Ehrenfeucht-Fraïssé games for FO In each move of the m-move Ehrenfeucht-Fraïssé game, the Spoiler picks an element in one of the two structures, and the Duplicator responds by picking an element of the other structure aiming to maintain a partial isomorphism. The m-move Ehrenfeucht-Fraïssé game is tailored for the fragment of FO consisting of all formulas of quantifier depth at most m. The family of the m-move Ehrenfeucht-Fraïssé games, $m \geq 1$, is the main tool for studying FO-definability on arbitrary class of structures and, in particular, on classes of finite structures. For instance, these games can be used to show that none of the following properties is FO-definable on the class $\mathcal{G}$ of all finite graphs: CONNECTIVITY, ACYCLICITY, PLANARITY, and 2-COLORABILITY.

Ehrenfeucht-Fraïssé games for ESO In these games, the Spoiler first picks a number of relations of specified arities on the structure $\mathbf{A}$ and the Duplicator responds by picking relations of matching arities on the structure $\mathbf{B}$. After this, the two players engage on a m-move Ehrenfeucht-Fraïssé game on the expanded structures obtained by adding the relations selected by the players. These games are parameterized by the number and the arities of the relations in the initial move and by the number m of the moves in the m-move Ehrenfeucht-Fraïssé game that follows.

Ehrenfeucht-Fraïssé games for ESO yield a sound and complete method for analyzing ESO-definability on classes of finite structures. Hence, in principle, if $\mathrm{NP} \neq \mathrm{coNP}$, then this separation can be proved using Ehrenfeucht-Fraïssé games for ESO by showing, for instance, that NON-3-COLORABILITY is not ESO-definable on the class $\mathcal{G}$ of all finite graphs. This methodology has been applied with considerable success in proving that certain interesting queries are not definable in *monadic* ESO, the fragment of existential second-order logic in which all existential second-order quantifiers range over unary relations (i.e., sets). For example, Fagin [29] used this method to show that CONNECTIVITY is not definable in monadic ESO on the class of all finite graphs, which implies that monadic ESO is not closed under complement in the finite. At the time, this gave a glimmer of hope that the method of combinatorial games could ultimately be used to establish the separation of NP from coNP. This early optimism, however, has yet to translate to breakthrough results in complexity.

One of the reasons for the lack of progress is that, although combinatorial games provide a sound and complete method for analyzing definability in many logics, the implementation of this method can be extremely difficult, since finding and spelling out the winning strategies for the Duplicator can be a combinatorially arduous task. In the case of first-order logic, this task becomes easier by taking advantage of the *locality* of FO, which, intuitively, is the property that a formula can only talk about "neighborhoods" of some fixed radius. Using this property, it is possible to give broad and easily checkable sufficient conditions for the Duplicator to win the game. Locality can also be used in the study of definability in monadic ESO [32]. (A detailed exposition of locality and its applications to definability in logics on finite structures can be found in [62].)

Locality arguments, however, cannot be used for fragments of ESO that are richer than monadic ESO; instead, insurmountable combinatorial obstacles seem to arise when attempting to apply the method of combinatorial games to such fragments of ESO. As a matter of fact, we know very little even about the expressive power of *binary* ESO, the fragment of ESO in which all existential second-order quantifiers range over binary relations. In particular, the following problem raised by Fagin [28] more than thirty years ago still remains open.

Problem 3 *Prove or disprove that there is a query Q on the class $\mathcal{G}$ of all finite graphs such that Q is* ESO*-definable (i.e., Q is in* NP*), but is not definable in binary* ESO*.*

It is not even known whether or not binary ESO formulas with a single existentially quantified binary symbol have the same expressive power as arbitrary ESO formulas on $\mathcal{G}$.

Pebble games for LFP and PFP In Section 3.1, we stated that the fixed-point logics LFP and PFP cannot express simple counting queries, such as EVEN CARDINALITY, on the class of all finite σ-structures. These limitations of the expressive power of LFP and PFP are proved by first viewing these logics as fragments of the *finite-variable infinitary logic* $\mathrm{L}^{\omega}_{\infty\omega}$ and then using combinatorial games for $\mathrm{L}^{\omega}_{\infty\omega}$.

$\mathrm{L}_{\infty\omega}$ is the extension of FO with infinite disjunctions and conjunctions. For every positive integer k, we let $\mathrm{L}^{k}_{\infty\omega}$ denote the collection of all $\mathrm{L}_{\infty\omega}$-formulas with at most k distinct variables (each variable, however, may be reused in a formula any number of times). Finally, by definition, $\mathrm{L}^{\omega}_{\infty\omega} = \bigcup_{k=1}^{\infty} \mathrm{L}^{k}_{\infty\omega}$. The finite-variable infinitary logic $\mathrm{L}^{\omega}_{\infty\omega}$ was introduced by Barwise [12] in the study of LFP on fixed infinite structures, but later on found numerous uses and applications in finite model theory.

It is known that the containments $\mathrm{LFP} \subseteq \mathrm{PFP} \subseteq \mathrm{L}^{\omega}_{\infty\omega}$ hold on the class $\mathcal{F}$ of all finite σ-structures [59]. Moreover, the containment of PFP in $\mathrm{L}^{\omega}_{\infty\omega}$ on $\mathcal{F}$ is a proper one, as $\mathrm{L}^{\omega}_{\infty\omega}$ can express non-recursive queries on $\mathcal{F}$. Barwise [12], and also Immerman [46], showed that, for every positive integer k, the expressive power of the k-variable infinitary logic $\mathrm{L}^{k}_{\infty\omega}$ can be analyzed using a variant of the Ehrenfeucht-Fraïssé games that has become known as k-*pebble* games. The k-pebble game is played on two structures $\mathbf{A}$ and $\mathbf{B}$. The Spoiler and the Duplicator have k pebbles each; they take turns and they place or remove at most k of their pebbles on or from elements of the two structures. The Spoiler wins if at some point of time the pebbled substructures are not isomorphic; the Duplicator wins if he can continue playing "forever" so that an isomorphism between the pebbled substructures is maintained.

The k-pebble game is the main tool for proving limitations in the expressive power of $\mathrm{L}^{k}_{\infty\omega}$, $k \geq 1$ and, a fortiori, in the expressive power of LFP and PFP. For example, since, for every $k \geq 1$, the Duplicator has a trivial winning strategy for the k-pebble game on the complete graphs $\mathbf{K}_k$ and $\mathbf{K}_{k+1}$ with k and $k+1$ nodes respectively, it follows that EVEN CARDINALITY is not $\mathrm{L}^{\omega}_{\infty\omega}$-definable on the class $\mathcal{G}$ of all finite graphs. Consequently, EVEN CARDINALITY is not LFP-definable (or PFP-definable) on $\mathcal{G}$. Since EVEN CARDINALITY is easily seen to be ESO-definable on $\mathcal{G}$, it follows that $\mathrm{LFP} \neq \mathrm{ESO} = \mathrm{NP}$ on $\mathcal{G}$. This separation, however, does imply that $\mathrm{P} \neq \mathrm{NP}$, since $\mathrm{LFP} = \mathrm{P}$ on the class of all ordered finite graphs, but not on the class $\mathcal{G}$ of all finite graphs.

The method of k-pebble games, $k \geq 1$, is sound and complete for analyzing definability in $\mathrm{L}^{\omega}_{\infty\omega}$. This method, however, does not help in distinguishing the expressive power of LFP from that of PFP, since these logics are properly contained in $\mathrm{L}^{\omega}_{\infty\omega}$ on finite structures. By Theorems 6 and 8, $\mathrm{LFP} = \mathrm{P}$ and $\mathrm{PFP} = \mathrm{PSPACE}$ on the class $\mathcal{O}$ of all ordered finite structures; hence, separating LFP from PFP on $\mathcal{O}$ amounts to separating P from PSPACE, yet another outstanding open problem in computational complexity. How do LFP and PFP compare on the class $\mathcal{F}$ of all finite σ-structures? This question was raised by Chandra and Harel [16], who posed the following problem.

Problem 4 (Chandra and Harel) *Show that* $\mathrm{LFP} \neq \mathrm{PFP}$ *on the class $\mathcal{F}$ of all finite σ-structures.*

No progress was made on this problem for almost a decade, until Abiteboul and Vianu [4] obtained the following unexpected result that explained the lack of progress.

Theorem 9 *The following statements are equivalent.*

- $\mathrm{LFP} \neq \mathrm{PFP}$ *on the class $\mathcal{F}$ of all finite σ-structures.*
- $\mathrm{P} \neq \mathrm{PSPACE}$.

As seen in Section 3.1, LFP is contained in P on $\mathcal{F}$, and PFP can express PSPACE-complete problems on $\mathcal{F}$; this implies that if $\mathrm{P} \neq \mathrm{PSPACE}$, then $\mathrm{LFP} \neq \mathrm{PFP}$ on $\mathcal{F}$. The proof of the other direction of Theorem 9 is much harder and uses deeper connections between LFP, k-pebble games, and $\mathrm{L}^k_{\infty\omega}$-types. Theorem 9 is one of the success stories of finite model theory; in particular, it reveals that the difference between polynomial-time computability and polynomial-space computability amounts to the difference between two different mechanisms for iterating first-order formulas on the class of all finite structures.

It remains a challenge to develop methods, other than the method of k-pebble games, to analyze the expressive power of LFP and PFP on finite structures. Note that the method of k-pebble games is of no use on classes of ordered finite structures, as every query on the class $\mathcal{O}$ of all ordered finite σ-structures is $\mathrm{L}^2_{\infty\omega}$-definable; the reason for this is that the isomorphism type of every ordered finite structure is definable by a first-order sentence with 2 distinct variables. As a matter of fact, even the difference between FO and LFP on classes of ordered finite structures is not well understood. In particular, the following conjecture, introduced in [58] and known as the *Ordered Conjecture*, remains open.

Conjecture 3 (Kolaitis and Vardi) *If $\mathcal{C}$ is a class of finite ordered structures of arbitrarily large finite cardinalities, then* $\mathrm{FO} \neq \mathrm{LFP}$ *on* $\mathcal{C}$.

It has been shown that either way of settling the Ordered Conjecture would resolve open problems in computational complexity [23, 24].

Existential pebble games and Datalog A *Datalog* program is a function-free and negation-free logic program. As a query language, Datalog has been investigated in depth by the database theory community in the context of deductive databases (see [2]). Datalog has the same expressive power as the *existential positive* fragment of LFP [17]. For example, the following Datalog program expresses the PATH SYSTEMS query (recall also that this query is P-complete):

$$\begin{array}{lcl} S(x) & :- & A(x) \\ S(x) & :- & R(x,y,z), S(y), S(z). \end{array}$$

The expressive power of Datalog can be analyzed using *existential k-pebble games*, which were introduced in [60] and shown to capture the expressive power of the existential fragment $\exists\mathrm{L}^\omega_{\infty\omega}$ of $\mathrm{L}^\omega_{\infty\omega}$. These games differ from the k-pebble games in two ways: first, the Spoiler plays always on the structure $\mathbf{A}$ and the Duplicator plays always on the structure $\mathbf{B}$; second, the goal of the Duplicator is to maintain a *homomorphism* (instead of an isomorphism) between the pebbled substructures. Extensions of Datalog with inequalities $\neq$ or with other limited forms of negation can also be analyzed using suitable variants of the existential k-pebble games. This is another manifestation of the adaptability and versatility of combinatorial games in studying the expressive power of logics.

3.3 Logic and Asymptotic Probabilities

Assume that $\mathcal{C}$ is a class of finite σ-structures and, for every $n \geq 1$, let $\mathcal{C}_n$ be the subclass of $\mathcal{C}$ consisting of all members of $\mathcal{C}$ with universe of size n. A *measure* on $\mathcal{C}$ is a sequence $\mu = \mu_n$, $n \geq 1$, of probability measures on $\mathcal{C}_n$. If Q is a Boolean query on $\mathcal{C}$, then we write $\mu_n(Q)$ to denote the probability of Q on $\mathcal{C}_n$. For example, if μ_n is the uniform probability measure on $\mathcal{C}_n$, then $\mu_n(Q)$ is equal to the fraction of the structures in $\mathcal{C}_n$ that satisfy Q. The *asymptotic probability* $\mu(Q)$ of the query Q is defined as $\mu(Q) = \lim_{n\to\infty} \mu_n(Q)$, provided the limit exists.

A great deal is known about the asymptotic probabilities of queries on finite graphs. As an example, for the uniform measure μ on finite graphs, it is known (and easy to see) that $\mu(\textsc{Connectivity}) = 1$ and $\mu(\text{2-}\textsc{Colorability}) = 0$. Note, though, that $\mu(\textsc{Even Cardinality})$ does not exist.

Let L be a logic, $\mathcal{C}$ a class of finite σ-structures, and μ a measure on $\mathcal{C}$. The *0–1 law holds for L on $\mathcal{C}$ w.r.t. μ* if $\mu(Q) = 0$ or $\mu(Q) = 1$, for every L-definable query Q on $\mathcal{C}$. The *convergence law holds for L on $\mathcal{C}$ w.r.t. μ* if $\mu(Q)$ exists for every L-definable query Q on $\mathcal{C}$.

The investigation of 0–1 laws in finite model theory started with a 0–1 law for first-order logic established by Glebskii et al. [34] and, independently, by Fagin [30].

Theorem 10 *The 0–1 law holds for FO on the class $\mathcal{F}$ of all finite σ-structures w.r.t. the uniform measure.*

Fagin [30] derived this 0–1 law by first proving the following result, which can be regarded as a *transfer theorem*.

Theorem 11 *Let μ be the uniform measure on the class $\mathcal{F}$ of all finite σ-structures. There is a unique countable graph* $\mathbf{R}$ *such that, for every* FO*-sentence ψ, the following two statements are equivalent:*

- $\mathbf{R} \models \psi$.
- $\mu(\psi) = 1$.

$\mathbf{R}$ is known as Rado's Graph, the unique countable homogeneous and *universal* graph (the latter means that $\mathbf{R}$ contains every finite graph as an induced subgraph). It is characterized by a set of first-order *extension axioms* that, intuitively, assert that every finite subgraph can be extended in every possible way. Theorem 11 is proved in two steps: first, a back and forth argument is used to show that the set of all extension axioms has a unique countable model; second, it is shown that the asymptotic probability of each extension axiom is equal to 1. Theorem 11 is a transfer principle between truth of FO-sentences on $\mathbf{R}$ and almost sure truth of FO-sentences on $\mathcal{F}$ w.r.t. the uniform measure. The 0–1 law for FO follows immediately from Theorem 11.

If the 0–1 law holds for a logic L, then there is a natural decision problem associated with it: given an L-sentence ψ, tell whether $\mu(\psi) = 0$ or $\mu(\psi) = 1$. Grandjean [38] pinpointed the computational complexity of the 0–1 law for FO on $\mathcal{F}$ w.r.t. the uniform measure by showing that this problem is PSPACE-complete. Thus, telling if an FO-sentence is *almost surely true* on all finite structures is a decidable problem (in fact, decidable in PSPACE), while, by Trakhtenbrot's Theorem, telling if an FO-sentence is *true* on all finite structures is an undecidable problem.

0–1 laws for extensions of FO There has been an extensive investigation of 0–1 laws for various extensions of FO on the class $\mathcal{F}$ of all finite σ-structures w.r.t. the uniform measure. The following result, obtained by Kolaitis and Vardi [59], subsumes earlier 0–1 laws for LFP [13] and for PFP [56]. The proof uses k-pebble games in a crucial way.

Theorem 12 *The 0–1 law holds for the finite-variable infinitary logic* $\mathrm{L}^{\omega}_{\infty\omega}$ *on the class* $\mathcal{F}$ *of all finite* σ*-structures w.r.t. the uniform measure.*

Note that 0–1 laws have implications for definability. In particular, Theorem 12 implies that EVEN CARDINALITY is not $\mathrm{L}^{\omega}_{\infty\omega}$-definable on $\mathcal{F}$; the same holds true for every query whose asymptotic probability does not exist or (exists and) is different from 0 and 1.

The convergence law (and, *a fortiori*, the 0–1 law) fails for ESO on $\mathcal{F}$ w.r.t. the uniform measure, since EVEN CARDINALITY is ESO-definable. In fact, the convergence law fails even for monadic ESO [51]. Nonetheless, many well-known NP-complete problems, including 3-COLORABILITY and SATISFIABILITY, have asymptotic probability 0 or 1. This motivated the study of 0–1 laws for fragments of ESO; one natural way to obtain such fragments is to use the quantifier prefix in the first-order part of ESO-formulas as a parameter. Specifically, if Ψ is a prefix class of FO-formulas, then $\mathrm{ESO}(\Psi)$ denotes the collection of all ESO-formulas of the form $\exists R_1 \cdots \exists R_k \varphi$ such that φ is a formula in Ψ. For example, it is not hard to see that 3-COLORABILITY is expressible by an $\mathrm{ESO}(\forall\forall)$-formula, while SATISFIABILITY is expressible by an $\mathrm{ESO}(\forall\exists)$-formula with CNF-formulas encoded by finite structures in a straightforward way.

The following theorem yields a complete classification, under the above parametrization, of the fragments of ESO for which the 0–1 law holds.

Theorem 13 *The Bernays-Schönfinkel Class* $\exists^*\forall^*$ *and the Ackerman Class* $\exists^*\forall\exists^*$ *are the only prefix classes* Ψ *of first-order logic such that the 0–1 law holds for the associated fragment* $\mathrm{ESO}(\Psi)$ *of* ESO *on the class* $\mathcal{F}$ *of all finite* σ*-structures w.r.t. the uniform measure.*

The key to establishing Theorem 13 was to first prove the following three results. The first two are due to Kolaitis and Vardi [56, 57], while the third one is due to Pacholski and Szwast [68].

Theorem 14 *Let* $\mathcal{F}$ *be the class of all finite* σ*-structures.*

- *The 0–1 law holds for the fragment* $\mathrm{ESO}(\exists^*\forall^*)$ *on* $\mathcal{F}$ *w.r.t. the uniform measure.*
- *The 0–1 law holds for the fragment* $\mathrm{ESO}(\exists^*\forall\exists^*)$ *on* $\mathcal{F}$ *w.r.t. the uniform measure.*
- *The convergence law fails for the fragment* $\mathrm{ESO}(\forall\forall\exists)$ *on* $\mathcal{F}$ *w.r.t. the uniform measure.*

The decision problem for the 0–1 law for $\mathrm{ESO}(\exists^*\forall^*)$ and for the 0–1 law for $\mathrm{ESO}(\exists^*\forall\exists^*)$ is NEXPTIME-complete; in contrast, the decision problem for the asymptotic probabilities of $\mathrm{ESO}(\forall\forall\exists)$-sentences is undecidable.

In Theorem 14, equality is allowed in the first two fragments for which the 0–1 law holds, and equality is used in showing that the 0–1 law fails for the third fragment. As stated in Section 2, the Bernays-Schönfinkel Class $\exists^*\forall^*$ and the Ackermann Class $\exists^*\forall\exists^*$ are the only two prefix classes of FO (with equality) for which the satisfiability problem is decidable. Thus, the classification of 0–1 laws for prefix fragments of ESO mirrors the classification of the satisfiability problem for prefix classes of FO (with equality). In general, however, the decidability of a fragment of FO need not imply that the 0–1 law holds for the associated fragment of ESO. Indeed, consider the Gödel Class without equality consisting of all equality-free $\exists^*\forall\forall\exists^*$ sentences; Gödel [35] showed that this class has the finite model property, hence the satisfiability problem for it is decidable. It had been conjectured that the 0–1 law holds for the associated ESO fragment; this conjecture was disproved by Le Bars [11].

Theorem 15 *The convergence law fails for the fragment* $\mathrm{ESO}(\forall\forall\exists)$ *without equality on the class* $\mathcal{F}$ *of all finite* σ*-structures w.r.t. the uniform measure.*

0–1 laws on restricted classes of structures A different line of investigation focused on 0–1 laws for FO on restricted classes of finite structures w.r.t. the uniform measure. This investigation entailed an extensive interaction of finite model theory with asymptotic combinatorics. In particular, Compton [19, 20] characterized the existence of a 0–1 law for FO on a class $\mathcal{C}$ of finite structures in terms of properties of the *exponential generating series* of $\mathcal{C}$, provided this series has a positive radius of convergence. This characterization, however, does not apply to "fast growing" classes, such as the class $\mathcal{F}$ of all finite σ-structures. For such classes, 0–1 laws for FO were obtained with different techniques. As a further illustration of the interaction of finite model theory with asymptotic combinatorics, we mention two results about fast growing classes.

Compton [18] showed that the 0–1 law holds for FO on the class of all finite partial orders w.r.t. the uniform measure. The proof of this result made use of a theorem by Kleitman and Rothschild [52] that describes the asymptotic structure of finite partial orders and, in particular, asserts that almost all finite partial orders have height exactly three. Kolaitis, Prömel, and Rothchild [55] showed that, for every $k \geq 3$, the 0–1 law holds for FO on the class of all $\mathbf{K}_k$-free graphs w.r.t. the uniform measure, where $\mathbf{K}_k$ is the complete graph with k nodes. This result was obtained after investigating the asymptotic structure of $\mathbf{K}_k$-free graphs and showing that, for every $k \geq 3$, almost all $\mathbf{K}_k$-free graphs are $(k-1)$-colorable.

0–1 laws for FO under variable measures The study of random graphs was initiated by Erdös and Rényi [27]. A random graph with n nodes is obtained by putting edges between two nodes with probability given by some function $p(n)$. If $p(n) = 1/2$, then the resulting measure is the uniform measure considered earlier. The most well studied variable probability measures are of the form $p(n) = n^{-\alpha}$, where α is a real number. Shelah and Spencer [73] investigated 0–1 laws for first-order logic FO on the class $\mathcal{G}$ of all finite graphs w.r.t. such non-uniform measures and obtained the following remarkable result. A detailed proof that combines Ehrenfeucht-Fraïssé games with techniques from random graphs can be found in [37, Chapter 4].

Theorem 16 *Let α be a real number between 0 and 1*

- *If α is irrational, then the 0–1 law holds for* FO *on $\mathcal{G}$ w.r.t. the measure $p(n) = n^{-\alpha}$.*
- *It α is rational, then the convergence law for* FO *fails on $\mathcal{G}$ w.r.t. the measure $p(n) = n^{-\alpha}$.*

It follows that if $\alpha \in (0,1)$ is rational, then the 0–1 law fails for $\mathrm{L}^{\omega}_{\infty\omega}$ on $\mathcal{G}$ w.r.t. the measure $p(n) = n^{-\alpha}$. McArthur [65] showed that the 0–1 law also fails for $\mathrm{L}^{\omega}_{\infty\omega}$ on $\mathcal{G}$ w.r.t. measures $p(n) = n^{-\alpha}$ when $\alpha \in (0,1)$ is irrational. Hence, the 0–1 law fails for $\mathrm{L}^{\omega}_{\infty\omega}$ on $\mathcal{G}$ w.r.t. every measure of the form $p(n) = n^{-\alpha}$, where $0 < \alpha < 1$.

We conclude this section with a brief assessment of the research in the area of 0–1 laws. On the positive side, 0–1 laws are new phenomena that that are special to finite model theory and significantly contribute to its distinct character. Here, finiteness is a feature, and not a limitation, as 0–1 laws are meaningful because only finite structures are considered. Furthermore, even the small sample of results highlighted here, makes it clear that the study of 0–1 laws has enhanced the interaction between logic and combinatorics in a way that has benefitted both communities. On the negative side, it is fair to say that this area of research had less contact with or impact on computer science than other areas of finite model theory did. There are, of course, the connections with computational complexity concerning the decision problem for 0–1 laws that we touched upon earlier. There are also interesting connections and interaction with random Boolean satisfiability [7, 8]. Nonetheless, in some of the early papers on 0–1 laws, one of the stated main motivations for pursuing this line of investigation was that the analysis of the asymptotic probabilities of logical properties may be useful in the average-case analysis of algorithms. While there has been some work on the average complexity of database queries definable in fixed-point logics (eg., [1]), the early optimism and expectation for impact on the average complexity of algorithms remain largely unrealized.

3.4 Classical model theory in the finite

Some classical results of model theory, such as the Löwenheim-Skolem Theorem, are not meaningful in the finite. Others, such as the Compactness Theorem for first-order logic, are easily seen to fail in the finite. Furthermore, using the fact that EVEN CARDINALITY is not FO-definable on $\mathcal{F}$, it is easy to show that also the Craig Interpolation Theorem fails in the finite. In regard to preservation theorems, however, the state of affairs, is more subtle.

As described in Section 2, Scott and Suppes [72] conjectured in 1958 that the Łoś-Tarski Theorem about preservation-under-substructures holds in the finite. Soon after the conjecture was made, Tait [74] disproved it by exhibiting an existential-universal ($\exists^*\forall^*$) FO-sentence that is preserved under substructures on all finite structures, but it is not equivalent to any universal FO-sentence on all finite structures. There was no in-depth investigation of preservation theorems in the finite until the 1980s, when the interest in them was rekindled. Gurevich's 1984 paper [42] examined the failure of classical results of model theory in the finite and also contained a proof of an unpublished result by Compton to the effect that Tait's counterexample was, in a certain sense, optimal. Specifically, Compton showed that if a universal-existential ($\forall^*\exists^*$) FO-sentence is preserved under substructures on all finite structures, then it is equivalent to some universal FO-sentence on all finite structures. After this, Ajtai and Gurevich [5] showed that yet another classical preservation theorem, the Lyndon Positivity Theorem, fails in the finite: there is an FO-sentence $\varphi(S)$ that is monotone in S on all finite structures, but is not equivalent to any positive-in-S FO-sentence on all finite structures.

By the mid to late 1980s, most classical preservation theorems were shown to fail in the finite. A notable exception was the *preservation-under-homomorphisms* theorem, which asserts that if a FO-sentence is preserved under homomorphisms on all structures, then it is logically equivalent to an existential positive FO-sentence. The question of whether this theorem holds in the finite was singled out as an important problem in finite model theory, and made the list of such open problems in finite model theory at http://www-mgi.informatik.rwth-aachen.de/FMT/; in addition, it attracted the attention of researchers in database theory, because existential positive sentences express *unions of conjunctive queries*, which are arguably the most-frequently-asked relational database queries. The status of the preservation-under-homomorphisms theorem in the finite remained in question until it was finally resolved by Rossman [70] in 2005.

Theorem 17 *Let $\mathcal{F}$ be the class of all finite σ-structures. If a* FO-*sentence ψ is preserved under homomorphisms on $\mathcal{F}$, then there is an existential positive FO-sentence ψ^* such that ψ is equivalent to ψ^* on $\mathcal{F}$.*

Preservation theorems do not *relativize*, i.e., if a preservation theorem holds for a class $\mathcal{C}$ of structures, then this theorem does not automatically hold for a subclass $\mathcal{C}'$ of $\mathcal{C}$. The reason is that if the statement of the theorem is restricted to a subclass $\mathcal{C}'$ of $\mathcal{C}$, then both the hypothesis and the conclusion of the theorem are weakened. Similarly, the failure of a preservation theorem does not relativize either.

Shortly before Theorem 17 was proved, Atserias, Dawar, and Kolaitis [10] investigated the status of the preservation-under-homomorphism theorem for restricted classes of finite structures, and showed that it holds for several such classes of interest in combinatorics and graph theory.

Theorem 18 *The following statements are true.*

- *Let $\mathcal{T}(k)$ be the class of all graphs of treewidth at most k, for some $k \geq 1$. If a* FO-*sentence is preserved under homomorphisms on $\mathcal{T}(k)$, then it is equivalent to an existential positive* FO-*sentence on $\mathcal{T}(k)$.*

- *If a* FO-*sentence is preserved under homomorphisms on all planar graphs, then it is equivalent to an existential positive* FO-*sentence on all planar graphs.*

Actually, both parts of Theorem 18 are special cases of a more general result, also proved in [10], about classes of structures whose Gaifman graph excludes at least one minor. It should be noted that these results are not implied by Theorem 17 (and are obtained with different methods).

In a related line of investigation, Atserias, Dawar, and Grohe [9] studied the Łoś-Tarski Theorem for restricted classes of finite structures and obtained a number of results, including the following.

Theorem 19 *The following statements are true.*

- *Let $\mathcal{T}(k)$ be the class of all graphs of treewidth at most k, for some $k \geq 1$. If a* FO-*sentence is preserved under substructures on $\mathcal{T}(k)$, then it is equivalent to a universal* FO-*sentence on $\mathcal{T}(k)$.*

- *There is a* FO-*sentence that is preserved under substructures on all planar graphs, but is not equivalent to any universal* FO-*sentence on all planar graphs.*

Theorems 7, 18, and 19 reveal that the classes $\mathcal{T}(k)$, $k \geq 1$, of graphs of bounded treewidth have good model-theoretic and descriptive-complexity-theoretic properties. More broadly, these results make a case that a lot can be gained by shifting the focus from the class of all finite structures to restricted classes of structures of combinatorial and graph-theoretic interest.

We conclude this section with a problem motivated from Lindström's characterization of first-order logic. Lindström [64] showed that first-order logic is a maximal logic possessing both the Compactness Theorem and the Löwenheim-Skolem Theorem. In other words, there is no logic that is strictly more expressive than first-order logic and possesses both the Compactness Theorem and the Löwenheim-Skolem Theorem. This result, obtained in 1969, became the catalyst for the development of *abstract model theory*.

Is there an *abstract finite model theory*? The following problem was posed in [54], and still remains open.

Problem 5 *Is there a Lindström-type characterization of first-order logic on finite structures? Similarly, is there a Lindström-type characterization of least fixed-point logic on finite structures?*

4 Concluding Remarks

As stated in the introduction, this paper is not a comprehensive survey of finite model theory. Inevitably, several topics in finite model theory received scant attention only or were not covered at all. For instance, we only glimpsed at the rich body of work on finite-variable logics, and we said even less about the work on logics with generalized quantifiers in the finite. Furthermore, we left out entirely the interaction of finite model theory with modal and temporal logics, which includes the connections with the μ-calculus and model checking. These are important topics that undoubtedly deserve their own separate overviews.

In regard to the interaction of model theory with other areas, we attempted to highlight some of the connections with computational complexity and with asymptotic combinatorics. There has been extensive and fruitful interaction with database theory, which, once again, we were able to cover at a superficial level only. Some of the earlier connections of finite model theory with database theory are detailed in the book [2]. Since that time, finite model theory has also found applications to constraint databases (see [62, Chapter 13] or [37, Chapter 5]) and has been influential in the foundations of information integration (see the surveys [53, 61]). A relatively new, and quite exciting, development is the increasing interaction between finite model theory and constraint satisfaction (see [37, Chapter 6]); it is interesting to note that this has been accompanied by the publication of several papers on constraint satisfaction in the proceedings of recent LICS symposia.

In conclusion, finite model theory has come a long way and has indeed evolved from a collection of early sporadic results to a mature research area. Finite model theory has had numerous successes in several different fronts, but there have been disappointments and frustrations along the way. In particular, the Spectrum Problem is still essentially open, the quest for a logic for P is still ongoing, and the early optimism about 0–1 laws having impact on the average-case analysis of algorithms has yet to materialize. It is fair to say, however, that the brick walls encountered by finite model theory (and, especially, by descriptive complexity) are not that much different from those encountered by areas of computational complexity, such as circuit complexity and propositional proof complexity. On the positive side, the shift of focus on restricted classes of finite structures is bearing fruit, and the growing connections with constraint satisfaction are promising. Taking everything into account, one can only hope that the next thirty years of finite model theory will be at least as fruitful as the past thirty.

Acknowledgment Many thanks to Albert Atserias, Ron Fagin, and Moshe Y. Vardi for their valuable comments on an earlier draft of this paper.

References

[1] S. Abiteboul, K. Compton, and V. Vianu. Queries are easier than you thought (probably). In *PODS*, pages 23–32, 1992.

[2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison Wesley, 1995.

[3] S. Abiteboul and V. Vianu. Fixpoint extensions of first-order logic and Datalog-like languages. In *LICS*, pages 71–79, 1989.

[4] S. Abiteboul and V. Vianu. Generic computation and its complexity. In *STOC*, pages 209–219, 1991.

[5] M. Ajtai and Y. Gurevich. Monotone versus positive. *Journal of the ACM*, 34:1004–1015, 1987.

[6] G. Asser. Das Repräsentantenproblem im Prädikatenkalkül der ersten Stufe mit identität. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 1:252–263, 1955.

[7] A. Atserias. Unsatisfiable random formulas are hard to certify. In *LICS*, pages 325–334, 2002.

[8] A. Atserias. Definability on a random 3-CNF formula. In *LICS*, pages 458–466, 2005.

[9] A. Atserias, A. Dawar, and M. Grohe. Preservation under extensions on well-behaved finite structures. In *ICALP*, pages 1437–1449, 2005.

[10] A. Atserias, A. Dawar, and Ph. G. Kolaitis. On preservation under homomorphisms and unions of conjunctive queries. *J. ACM*, 53(2):208–237, 2006; prelim. version in *PODS 2004*.

[11] J.-M. L. Bars. Fragments of existential second-order logic without 0-1 laws. In *LICS*, pages 525–536, 1998.

[12] J. Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42:292–296, 1977.

[13] A. Blass, Y. Gurevich, and D. Kozen. A zero–one law for logic with a fixed point operator. *Information and Control*, 67:70–90, 1985.

[14] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Springer, 1997.

[15] J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12:389–410, 1992.

[16] A. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences*, 25:99–128, 1982.

[17] A. Chandra and D. Harel. Horn clause queries and generalizations. *Journal of Logic Programming*, 1:1–15, 1985.

[18] K. Compton. The computational complexity of asymptotic problems i: Partial orders. *Inf. Comput.*, 78(2):108–123, 1988.

[19] K. J. Compton. A logical approach to asymptotic combinatorics I: First-order properties. *Advances in Mathematics*, 65:65–96, 1987.

[20] K. J. Compton. A logical approach to asymptotic combinatorics II: Monadic second-order properties. *J. Combin. Theory Ser. A*, 50:110–131, 1989.

[21] S. A. Cook. An observation of time-storage trade-off. *Journal of Computer and System Sciences*, 9:308–316, 1974.

[22] A. Dawar. Generalized quantifiers and logical reducibilities. *Journal of Logic and Computation*, 5(2):213–226, 1995.

[23] A. Dawar and L. Hella. The expressive power of finitely many generalized quantifiers. In *LICS*, pages 20–29, 1994.

[24] A. Dawar, S. Lindell, and S. Weinstein. First order logic, fixed point logic and linear order. In *CSL*, pages 161–177, 1995.

[25] R. Diestel. *Graph Theory*. Springer, 2nd edition, 2000.

[26] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.

[27] P. Erdös and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 7:17–61, 1960.

[28] R. Fagin. Generalized first–order spectra and polynomial–time recognizable sets. In R. M. Karp, editor, *Complexity of Computation, SIAM-AMS Proceedings, Vol. 7*, pages 43–73, 1974.

[29] R. Fagin. Monadic generalized spectra. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:89–96, 1975.

[30] R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41:50–58, 1976.
[31] R. Fagin. Finite-model theory—a personal perspective. *Theoretical Computer Science*, 116(1):3–31, 1993.
[32] R. Fagin, L. Stockmeyer, and M. Y. Vardi. On monadic NP vs. monadic co-NP. *Information and Computation*, 120(1):78–92, July 1995.
[33] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
[34] 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*, 5:142–154, 1969.
[35] K. Gödel. Ein spezialfall des entscheidungsproblem der theoretischen logik. *Ergebn. math. Kolloq.*, 2:27–28, 1932.
[36] W. D. Goldfarb. The Gödel class with equality is unsolvable. *Bull. Amer. Math. Soc. (New Series)*, 10:113–115, 1984.
[37] E. Grädel, Ph. G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Vardi, Y. Venema, and S. Weinstein. *Finite Model Theory and its Applications*. Springer, 2007.
[38] E. Grandjean. Complexity of the first–order theory of almost all structures. *Information and Control*, 52:180–204, 1983.
[39] M. Grohe. Complete problems for fixed–point logics. *Journal of Symbolic Logic*, 60(2):517–527, 1995.
[40] M. Grohe. Fixed-point logics on planar graphs. In *LICS*, pages 6–15, 1998.
[41] M. Grohe and J. Mariño. Definability and descriptive complexity on databases of bounded tree-width. In *ICDT*, pages 70–82, 1999.
[42] Y. Gurevich. Toward logic tailored for computational complexity. In M. M. Richter et al., editor, *Computation and Proof Theory, Lecture Notes in Mathematics 1104*, pages 175–216. Springer-Verlag, 1984.
[43] Y. Gurevich. Logic and the challenge of computer science. In E. Börger, editor, *Current trends in theoretical computer science*, pages 1–57. Computer Science Press, 1988.
[44] Y. Gurevich and S. Shelah. Fixed-point extensions of first-order logic. *Annals of Pure and Applied Logic*, 32:265–280, 1986.
[45] L. Hella. Logical hierarchies in PTIME. *Information and Computation*, 129(1):1–19, 1996; prelim. version in *LICS 1992*.
[46] N. Immerman. Upper and lower bounds for first-order expressibility. *Journal of Computer and System Sciences*, 25:76–98, 1982.
[47] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.
[48] N. Immerman. *Descriptive Complexity*. Springer, 1999.
[49] N. Immerman and E. S. Lander. Describing graphs: a first-order approach to graph canonization. In A. Selman, editor, *Complexity Theory Retrospective*, pages 59–81. Springer-Verlag, 1990.
[50] N. G. Jones and A. Selman. Turing machines and the spectra of first-order formulas. *Journal of Symbolic Logic*, 39:139–150, 1974.
[51] M. Kaufmann. Counterexample to the 0-1 law for existential monadic second-order logic. CLI Internal Note 32, Computational Logic Inc., December 1987.
[52] D. Kleitman and B. Rothschild. Asymptotic enumeration of partial orders on a finite set. *Trans. Amer. Math. Soc.*, 205:205, 1975.
[53] Ph. G. Kolaitis. Schema mappings, data exchange, and metadata management. In *PODS*, pages 61–75, 2005.
[54] Ph. G. Kolaitis and J. A. Väänänen. Generalized quantifiers and pebble games on finite structures. *Ann. Pure Appl. Logic*, 74(1):23–75, 1995; prelim. version in LICS 1992.
[55] Ph. G. Kolaitis, H. J. Prömel, and B. L. Rothschild. K_{l+1}-free graphs: asymptotic structure and a 0-1 law. *Trans. Amer. Math. Soc.*, 303:637–671, 1987.
[56] Ph. G. Kolaitis and M. Y. Vardi. The decision problem for the probabilities of higher-order properties. In *STOC*, pages 425–435, 1987.
[57] Ph. G. Kolaitis and M. Y. Vardi. 0-1 laws and decision problems for fragments of second-order logic. *Information and Computation*, 87:302–338, 1990.
[58] Ph. G. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *LICS*, pages 46–57, 1992.
[59] Ph. G. Kolaitis and M. Y. Vardi. Infinitary logic and 0-1 laws. *Information and Computation*, 98:258–294, 1992; prelim. version in *LICS 1990*. Logic in Computer Science.
[60] Ph. G. Kolaitis and M. Y. Vardi. On the expressive power of Datalog: tools and a case study. *Journal of Computer and System Sciences*, 51(1):110–134, August 1995; prelim. version in *PODS 1990*.
[61] M. Lenzerini. Data integration: A theoretical perspective. In *PODS*, pages 233–246, 2002.
[62] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
[63] S. Lindell. An analysis of fixed point queries on binary trees. *Theoretical Computer Science*, 85:75–95, 1991.
[64] P. Lindström. On extensions of elementary logic. *Theoria*, 35:1–11, 1969.
[65] M. McArthur. Convergence and 0-1 laws for $L^k_{\infty\omega}$ under arbitrary measures. In *CSL*, pages 228–241, 1994.
[66] J. Nurmonen. Counting modulo quantifiers on finite linearly ordered trees. In *LICS*, pages 484–493, 1996.
[67] M. Otto. *Bounded Variable Logics and Counting: A Study in Finite Models*. Springer, 1997.
[68] L. Pacholski and W. Szwast. A counterexample to the 0-1 law for the class of existential second-order minimal Gödel sentences with equality. *Information and Control*, 107:91–103, 1993; prelim. version in *LICS 1991*.
[69] B. Rossman. Successor-invariance in the finite. In *LICS*, pages 148–157, 2003.
[70] B. Rossman. Existential positive types and preservation under homomorphisisms. In *LICS*, pages 467–476, 2005.
[71] H. Scholz. Problem #1: Ein ungelóstes problem in der symbolischen logik. *Journal of Symbolic Logic*, 17:160, 1952.
[72] D. Scott and P. Suppes. Foundational aspects of theories of measurement. *Journal of Symbolic Logic*, 23(2):113–128, 1958.
[73] S. Shelah and J. Spencer. Zero-one laws for sparse random graphs. *J. Amer. Math. Soc.*, 1:97–115, 1988.
[74] W. W. Tait. A counterexample to a conjecture of Scott and Suppes. *Journal of Symbolic Logic*, 24:15–16, 1959.
[75] B. A. Trakhtenbrot. The impossibilty of an algorithm for the decision problem for finite models. *Doklady Akademii Nauk SSR*, 70:569–572, 1950.
[76] M. Y. Vardi. The complexity of relational query languages. In *STOC*, pages 137–146, 1982.