Random Sparse Unary Predicates

Saharon Shelah
Mathematics Department, Hebrew University, Jerusalem, Israel

Joel Spencer
Courant Institute of Mathematical Sciences, New York University, New York, New York 10012

ABSTRACT

Random unary predicates U on $[n]$ holding with probability p are examined with respect to sentences A in a first-order language containing U and "less than." When $p = p(n)$ satisfies $np^{k+1} \ll 1 \ll np^k$ it is shown that $\Pr[A]$ approaches a limit dependent only on k and A. In a similar circular model the limit is shown to be zero or one. © 1994 John Wiley & Sons, Inc.

1. INTRODUCTION

Let n be a positive integer, $0 \leq p \leq 1$. The random unary predicate $U_{n,p}$ is a probability space over predicates U on $[n] = \{1, \ldots, n\}$ with the probabilities determined by

$$\Pr[U(x)] = p\,, \quad 1 \leq x \leq n\,,$$

and the events $U(x)$ being mutually independent over $1 \leq x \leq n$. Informally, we think of flipping a coin for each x to determine if $U(x)$ holds, the coin coming up "heads" with probability p. We shall examine the first order language with equality, a unary predicate U and a binary predicate $\leq$. Examples of sentences in this language are:

$$A: \exists_x U(x)$$

$$B: \exists_x U(x) \wedge \forall_y \neg y < x$$

$$C: \exists_{x,y} U(x) \wedge U(y) \wedge \forall_z \neg [x < z \wedge z < y]$$

Random Structures and Algorithms, Vol. 5, No. 3 (1994)
© 1994 John Wiley & Sons, Inc. CCC 1042-9832/94/030375-20

Sh:432

($>$, $\geq$, $<$ are naturally definable from $\leq$ and equality.) For any such sentence S we have the probability

$$\Pr[U_{n,p} \models S]$$

While the use of unary predicates is natural for logicians, there are two other equivalent formulations that will prove useful. We may think of U as a subset of $[n]$ and speak about $i \in U$ rather than $U(i)$. Second we may associate with U a sequence of zeroes and ones where the ith term is one if $U(i)$ and zero if $\neg U(i)$. Thus we may talk of starting at i and going to the next one. We shall use all three formulations interchangeably.

Ehrenfeucht [2] showed that for any constant p and any sentence S in this language

$$\lim_{n\to\infty} \Pr[U_{n,p} \models S]$$

exists. In the case of sentences A and C the limiting probability is 1 whenever $p > 0$. But sentence B effectively states $1 \in U$; hence its limiting probability is p. We get around these edge effects with a new language, consisting of equality, a unary predicate U, and a ternary predicate C. We consider C as a built in predicate on $[n]$ with $C(x, y, z)$ holding if and only if either $x < y < z$ or $y < z < x$ or $z < x < y$. Thinking of $[n]$ as a cycle, with 1 coming directly after n, $C(x, y, z)$ is the relation that x to y to z goes in a clockwise direction. For any sentence S in this new language we can again define $\Pr[U_{n,p} \models S]$ only in this case Ehrenfeucht's results give a Zero–One Law: For any constant p and sentence S

$$\lim_{n\to\infty} \Pr[U_{n,p} \models S] = 0 \text{ or } 1.$$

We shall call the first language the *linear* language and the second language the *circular* language. As a general guide, the circular language will tend to Zero–One Laws while the linear language, because of edge effects, will tend to limit laws.

We shall not restrict ourselves to p constant but rather consider $p = p(n)$ as a function of n. We have in mind the "Evolution of Random Graphs" as first developed by Erdős and Rényi. Here as $p = p(n)$ evolves from zero to one, the unary predicate evolves from holding for no x to holding for all x. Analogously (but without formal definition) we have *threshold functions* for various properties. For example, $p(n) = n^{-1}$ is a threshold property for A. When $p(n) \ll n^{-1}$, almost surely A fails while when $p(n) \gg n^{-1}$ almost surely A holds. In [4] we showed that when $p = n^{-\alpha}$ with $\alpha \in (0, 1)$, irrational then a Zero–One Law held for the random graph $G(n, p)$ and in [3] we found a near characterization of those $p = p(n)$ for which the Zero–One Law held. The situation with random unary predicates turns out to be somewhat simpler.

Definition. *$p = p(n)$ satisfies the Zero–One Law for circular unary predicates if for every sentence S in the circular language*

$$\lim_{n\to\infty} \Pr[U_{n,p(n)} \models S] = 0 \text{ or } 1 .$$

Here is our main result.

Theorem 1.1. *Let k be an arbitrary positive integer and let $p = p(n)$ be such that $np^k \to \infty$ and $np^{k+1} \to 0$. Then $p = p(n)$ satisfies the Zero–One Law for circular unary predicates.*

The proof will require the preliminaries of Sections 2 and 3 and is given in Section 4. This is the only difficult part in achieving the following full characterization.

Theorem 1.2. *Let $p = p(n)$ be such that $p(n) \in [0, 1]$ for all n and either*

$$p(n) \ll n^{-1}$$

or for some positive integer k

$$n^{-1/k} \ll p(n) \ll n^{-1/(k+1)}$$

or for all $\epsilon > 0$

$$n^{-\epsilon} \ll p(n) \text{ and } n^{-\epsilon} \ll 1 - p(n)$$

or for some positive integer k

$$n^{-1/k} \ll 1 - p(n) \ll n^{-1/(k+1)}$$

or

$$1 - p(n) \ll n^{-1} .$$

Then $p(n)$ satisfies the Zero–One Law for circular unary predicates. Inversely, if $p(n)$ falls into none of the above categories, then it does not satisfy the Zero–One Law for circular unary predicates.

Proof (*assuming Theorem* 1.1). The inverse part is relatively simple. Let A_k be the sentence that there exist k consecutive elements $x_1, \ldots, x_k \in U$. [x, y are consecutive if for no z is $C(x, z, y)$. For $k = 2$ this is example C.] Then $\Pr[A_k]$ is (for a given n) a monotone function of p. When $p(n) \sim cn^{-1/k}$ and c is a positive constant, the probability $\Pr[A_k]$ approaches a limit strictly between zero and one. (Roughly speaking, $n^{-1/k}$ is a threshold function for A_k.) Thus for $p(n)$ to satisfy the Zero–One Law, we must have $p(n) \ll n^{-1/k}$ or $p(n) \gg n^{-1/k}$. Further (replacing U with $\neg U$), the same holds with $p(n)$ replaced by $1 - p(n)$. For $p(n)$ to fall between these cracks, it must be in one of the above five categories.

When $p(n) \ll n^{-1}$, the Zero–One Law is trivially satisfied since almost surely there is no x for which $U(x)$. Also, if $p(n)$ satisfies the Zero–One Law, so does $1 - p(n)$. Suppose $p = p(n)$ satisfies $p(n) \gg n^{-\epsilon}$ and $1 - p(n) \gg n^{-\epsilon}$ for all $\epsilon > 0$. Theorem 2.10 gives that for every t there is a sequence $A_1 \cdots A_R$ with the property that for any sentence A of quantifier depth t either all models $\langle [u], C, U \rangle$ that contain $A_1 \cdots A_R$ as a consecutive subsequence satisfy A or no such models satisfy A. For $p(n)$ in this range $\langle [u], C, U \rangle$ almost surely contains

Sh:432

any such fixed sequence $A_1 \cdots A_R$ as a consecutive subsequence and hence the Zero–One Law is satisfied.

This leaves only the case of Theorem 1.1. □

Remark. Dolan [1] has shown that $p(n)$ satisfies the Zero–One Law for *linear* unary predicates if and only if $p(n) \ll n^{-1}$ or $n^{-1} \ll p(n) \ll n^{-1/2}$ or $1 - p(n) \ll n^{-1}$ or $n^{-1} \ll 1 - p(n) \ll n^{-1/2}$. For $n^{-1/2} \ll p(n) = o(1)$, he considered the following property:

$$D\colon\ \exists_x U(x) \wedge [U(x+1) \vee U(x+2)] \\ \wedge \neg\exists_y [U(y) \wedge [U(y+1) \vee U(y+2)] \wedge y < x] \wedge U(x+1)\,.$$

(Addition is *not* in our language but we write $x + 1$ as shorthand for that z for which $x < z$ but there is no w with $x < w < z$.) In our zero-one formulation, D basically states that the first time we have 11 comes before the first time we have 101. This actually has limiting probability 0.5. This example illustrates that limiting probability for linear unary predicates can depend on edge effects and not just edge effects looking at U on a fixed size set $1, \ldots, k$ or $n, n-1, \ldots, n-k$.

The basic aim of this paper is to give a proof of Theorem 1.1. The results of Section 2 are background results from logic and, in the last section, probability and may be skimmed or used as a reference. The central argument appears in Section 3 in which an infinite model appears, Theorem 3.2 being a crucial step. The final steps of the proof of Theorem 1.1 are given in Section 4. The linear case is dealt with in Section 5.

2. BACKGROUND

2.1. Sequences

Let A be a fixed finite alphabet. In applications A will be $\{0\}$ or $\{0,1\}$ or P_k. However, in what follows only the size of A matters. We shall take our examples with $A = \{a, b, c\}$. Let ΣA denote the space of finite sequences $a_1 \cdots a_u$ of elements of A. We include the null sequence, denoted O. We associate with each sequence a model $\langle [u], \leq, f \rangle$ where $f\colon [u] \to A$ is given by $f(i) = a_i$.

Consider the first order language with equality and $\leq$ and with function symbol f and basic primitives $f(x) = \alpha$ for each variable x and each $\alpha \in A$. A typical sentence would be

$$\exists_x \exists_y [f(x) = a \wedge f(y) = b \wedge \forall_z \neg[x < y \wedge y < z]\,,$$

with the meaning that ab occurs in the sequence as consecutive terms. This naturally generalizes the language of the introduction, with $A = \{0, 1\}$. We call this the linear language for A.

We shall use four facts about these models. First we fix a positive integer t.

Definition. *Two models M_1, M_2 are equivalent if they have the same truth value on all first order sentences of quantifier depth at most t.*

Property A. *There are only a* finite *number of equivalence classes.*

Definition. *Let M denote the set of equivalence classes. If a model M_1 belongs to class $m_1 \in M$, we call m_1 the* Ehrenfeucht value *of M_1.*

Remark. The number of such classes may be very large. With A fixed, the number grows like a tower function as a function of t. While this makes calculation very difficult in our work, t is fixed so that M is of fixed size, and we do not concern ourself with its actual size.

We define addition of models by concatenation. For example, $aba + acca = abaacca$. More formally:

Definition. *Let $M_1 = \langle [u], \leq, f \rangle$, $M_2 = \langle [v], \leq, g \rangle$. We define $M_1 + M_2 = \langle [u + v], \leq, h \rangle$ where $h(i) = f(i)$ for $1 \leq i \leq u$ and $h(i) = g(i-u)$ for $u < i \leq u+v$.*

Property B. *If M_1, M_1' are equivalent and M_2, M_2' are equivalent, then $M_1 + M_2$ is equivalent to $M_1' + M_2'$.*

Definition. *Let $m_1, m_2 \in M$. Let M_1, M_2 be any models having m_1, m_2 respectively as their Ehrenfeucht values. We define $m_1 + m_2$ to be the Ehrenfeucht value of $M_1 + M_2$.*

The definition of $m_1 + m_2$ is unique by Property B. As addition on models is clearly associative, so is addition on M so that M forms a semigroup.

Notations. We let O denote both the null sequence and its Ehrenfeucht value, depending on context. Note that in M we have $m + O = O + m = m$ for all m. For each $a \in A$ we let a also denote the Ehrenfeucht value of the sequence a, i.e., the model $\langle [1], \leq, f \rangle$ with $f(1) = a$.

The third basic fact concerns copies of a given model. For any positive integer r and any $m \in M$, we define rm as the sum of r m's. More formally, by induction, $1m = m$ and $(r+1)m = rm + m$. We set $s = 3$ for definiteness below, though that is not precisely the best value.

Property C. *For all $m \in M$ and all $i, j > s$*

$$im = jm .$$

The final fact tells us that the Ehrenfeucht values truly reflect the world of first order sentences of quantifier depth t. We can correspond any first order sentence B of quantifier depth at most t with the set $S \subseteq M$ of equivalence classes of models satisfying B.

Property D. For every $S \subseteq M$ there is a first order sentence B of quantifier depth at most t such that B is satisfied by precisely those models whose Ehrenfeucht value lies in S.

Sh:432

2.2. Persistent and Transient

The following theorems hold in any finite semigroup M with identity and with Property C. Their application to ΣA is deferred to the next section.

Theorem 2.1 (and Definition). *We call $x \in M$* persistent *if*

$$\forall_y \exists_z x + y + z = x , \tag{1}$$

$$\forall_y \exists_z z + y + x = x , \tag{2}$$

$$\exists_p \exists_s \forall_y p + y + s = x . \tag{3}$$

These three properties are equivalent. We call x transient *if it is not persistent.*

Proof of Equivalence.
(3)$\Rightarrow$(1): Take $z = s$, regardless of y. Then

$$x + y + z = (p + y + s) + y + s = p + (y + s + y) + s = x .$$

(1)$\Rightarrow$(3): Let $R_x = \{x + v:\ v \in M\}$. We first claim there exists $u \in M$ with $|R_x + u| = 1$, i.e., $x + y + u$ remains the same for all y. Otherwise take $u \in M$ with $|R_x + u|$ minimal and say $v, w \in R_x + u$. As $R_x + u \subseteq R_x$, we write $v = x + u_1$, $w = x + u_2$. From (1), with $y = u_1$, we have $x = v + u_3$ and thus $w = v + u_4$ with $u_4 = u_3 + u_2$. From Property C there is an integer q with $qu_4 = (q + 1)u_4$. Then

$$w + qu_4 = v + (q + 1)u_4 = v + qu_4 .$$

Adding qu_4 to $R + u$ sends v, w to the same element so $|R + u + qu_4| < |R + u|$, contradicting the minimality. Now say $R_x + u = \{u_5\}$. Again by (1) there exists u_6 with $u_5 + u_6 = x$. Then $R_x + (u + u_6) = \{x\}$ so that (3) holds with $p = x$, $s = u + u_6$.

By reversing addition [noting that (3) is self-dual while the dual of (1) is (2)], these arguments give that (3) and (2) are equivalent. □

Remark. Define a directed graph G^r on M by directing edges from x to $x + y$ for every $x, y \in M$. From (1) the persistent elements of M are precisely the elements of the strongly connected components of G^r.

Theorem 2.2. *If z is persistent and $c \in M$, then*

$$z + c + z = z .$$

Proof. Let p, s, as given by (3), have the property that $p + y + s = z$ for all y. Taking $y = O$, $p + s = z$. Then $z + c + z = p + (s + c + p) + s = z$.

Theorem 2.3. *If x is persistent, then $w_1 + x + w_2$ is persistent for all $w_1, w_2 \in M$.*

Proof. Let x be persistent and set $v = w_1 + x + w_2$. For any $y \in M$ set $z = v$ so that

$$v + y + z = w_1 + [x + (w_2 + y + w_1) + x] + w_2 = w_1 + x + w_2 = v\ ,$$

and hence v is persistent. □

Definition. *We define the relation* $x \equiv_R u$ *by* $\exists_v(x + v = u)$. *We define the relation* $x \equiv_L u$ *by* $\exists_v(v + x = u)$.

Theorem 2.4. $\equiv_R$ *and* $\equiv_L$ *are equivalence relations on the set of persistent elements of* M.

Proof. Immediate from (1) and (2), respectively. □

Definition. *We define*

$$R_x = \{x + v\colon v \in M\}\ ,$$
$$L_x = \{v + x\colon v \in M\}\ .$$

Remark. For x persistent, R_x is the strongly connected component of G^r containing x.

Theorem 2.5. *For* x *persistent,* R_x *is the equivalence class containing* x *under* $\equiv_R$ *and* L_x *is the equivalence class containing* x *under* $\equiv_L$.

Proof. Immediate.

Theorem 2.6. *Let* x, y *be persistent. Then*

$$R_x \cap L_y = \{x + y\}\ .$$

Proof. Clearly $x + y \in R_x \cap L_y$. Let $z \in R_x \cap L_y$. Then there exist a, b with $x = z + a$ and $y = b + z$ so that $x + y = z + (a + b) + z$. But z is persistent, and by Theorem 2.2 $z + (a + b) + z = z$, so that $x + y = z$. □

One final theorem shows that these ideas are not pointless.

Theorem 2.7. *There exists a persistent* $x \in M$.

Proof. Take x with $x + M$ of minimal size. For any $y \in M$, $x + y + M \subseteq x + M$ so that $x + y + M = x + M$. As $x = x + O \in x + M$, $x \in x + y + M$ so that there exists z with $x + y + z = x$, and x is persistent. □

2.3. Persistent Sequences

We fix a finite alphabet A and the parameter t of the previous section. In examples, we shall consider $A = \{a, b, c\}$, $t = 5$, and $s = 243$ satisfying Property C of the previous section.

Sh:432

Definition. *A sequence $\sigma = a_1 \cdots a_r \in \Sigma A$ is called persistent if its Ehrenfeucht value m is persistent in the sense of the previous section. Otherwise, we call σ transient.*

Our object in this section is to get a reasonable picture of what a persistent sequence is. Our first result is a sufficiency condition.

Theorem 2.8. *There is a $\sigma \in \Sigma A$ so that any $\tau \in \Sigma A$ containing σ as a consecutive subsequence is persistent.*

Proof. Pick any σ whose Ehrenfeucht value m is persistent. Any $\tau = \gamma^- + \sigma + \gamma^+$ has Ehrenfeucht value $m^- + m + m^+$ which is persistent by Theorem 2.3.

The following definition and theorem are not formally necessary in our presentation, but we feel they give a good clue as to what persistency really means.

Definition. *A first order sentence B (in the linear language for A) is called* central *if (i) some $\sigma \in \Sigma A$ satisfies B and (ii) if σ satisfies B and τ contains σ as a consecutive subsequence, then τ satisfies B.*

Theorem 2.9. *σ is persistent if and only if it satisfies all central sentences B of quantifier depth at most t.*

Proof. Assume σ persistent and B central. Let σ_0 satisfy B. Then $\tau = \sigma + \sigma_0 + \sigma$ satisfies B. But τ has the same Ehrenfeucht value as σ so they have the same truth value on B and so σ satisfies B.

For the converse we use Property D. Let B be the sentence corresponding to the set S of persistent states. From Theorem 2.3, B is a central property. If σ satisfies it, then the Ehrenfeucht value of σ, and hence σ itself, must be persistent. □

Roughly, the central sentences are existential statements that do not depend on the "edges" of a sequence. Let B be the sentence that the sequence begins with a, formally $\exists_x f(x) = a \wedge \neg\exists_y y < x$. This sentence is not central: If you take a sequence starting with a and add a b on the left, then it no longer has this property. Similar noncentral properties would be that the first non-a is a c or that the first time either aba or aca occurs as a consecutive subsequence, it is aba. All of these can have their truth value changed by changing the edges of the sequence. A typical central statement is that $accab$ appears as a subsequence. Once it does, no additions to the sequence on the edges can make it false. A more complicated statement is that there exists two a's with only b's between them:

$$\exists_{x,y} f(x) = a \wedge f(y) = a \wedge x < y \wedge [\forall_z (x < z \wedge z < y) \Rightarrow f(z) = b]\,.$$

Again, once this appears it cannot be destroyed by adding to the sequence at the edges. Thus a persistent sequence σ (when $t \geq 3$, the quantifier depth of this sentence) *must* have two a's with only b's between them.

2.4. Circular Sequences

Let A be a fixed finite alphabet. Let $Cyc(A)$ denote the space of finite sequences $a_1 \cdots a_u$ of elements of A, so that formally $Cyc(A) = \Sigma A$. We associate with each sequence a model $\langle [u], C, f \rangle$ where f: $[n] \to A$ is given by $f(i) = a_i$ and C is the built in "clockwise" ternary relation: $C(x, y, z)$ if and only if $x < y < z$ or $y < z < x$ or $z < x < y$. The circular language (for A) is the first order language with equality and C and with function symbol f and basic primitives $f(x) = \alpha$ for each variable x and each $\alpha \in A$. We think of the sequences as lying in a circle. There is a natural notion of σ being a consecutive subsequence of τ. (Formally, $a_1 \cdots a_u$ is a consecutive subsequence of $b_1 \cdots b_l$ if $u \le l$ and there exists s so that $b_{s+^*i} = a_i$, $1 \le i \le u$, where $s + {}^*i$ is $s + i$ when $s + i \le n$ and otherwise $s + i - n$.) Our main result can be thought of as stating the existence of a universal sequence. We first need compare the circular and linear modes.

Given a circular model $M = \langle [u], C, f \rangle$, we can naturally cut it at any $i \in [u]$ giving a linear model M_i. Informally, if M is the sequence $a_1 \cdots a_u$ [considered as a member of $Cyc(a)$], then M_i is the sequence $a_i \cdots a_u \cdot a_1 \cdots a_{i-1}$. Let $P(x)$ be a formula in the circular language. Consider all uses of the ternary relation in the formula P. When $C(y_1, y_2, y_3)$ does not have variable x replace it with $C(y_1, y_2, y_3)$ by $y_1 < y_2 < y_3 \vee y_2 < y_3 < y_1 \vee y_3 < y_1 < y_2$. Replace $C(x, y_1, y_2)$ with $x < y_1 < y_2$ and $C(y_1, x, y_2)$ with $x < y_2 < y_1$ and $C(y_1, y_2, x)$ by $x < y_1 < y_2$. Call the resulting formula $P^*(x)$; it is a formula in the linear language. The formula $\neg\exists_w w < x$ has the meaning that x is "first" in the linear language. This allows us to unravel the circular language. *The formula $P(x)$ is satisfied in M with $x = i$ if and only if $\exists_x[[\neg\exists_w w < x] \wedge P^*(x)]$ is satisfied by M_i.*

Lemma. *The truth value of all sentences up to quantifier depth t for a circular model $M = \langle [u], C, f \rangle$ is determined by the set of Ehrenfeucht value at level t of the linear models M_i, $1 \le i \le u$.*

Proof. Any such sentence can be expressed as the purely Boolean combination of sentences of the form $\exists_x P(x)$ and these are true if and only if some M_i satisfies $\exists_x[[\neg\exists_w w < x] \wedge P^*(x)$, which is of the same quantifier depth. □

Theorem 2.10. *For every t, A there exists a sequence $\sigma \in Cyc(A)$ so that all $\tau \in Cyc(A)$ that contain σ as a consecutive subsequence have the same truth values on all sentences (in the circular language) of quantifier depth at most t.*

Proof. For every persistent $x \in M$, let X be the specific sequence with this value, and let X^r denote the sequence X in reverse order. Let σ be the concatenation of all the sequences $X^r + X$. Consider the level t Ehrenfeucht values of the linear M_i. Any i cuts at most one of the X so that M_i will contain some (in fact, many) persistent X as a consecutive subsequence and hence, by Theorem 2.3, will be persistent. Conversely, given any persistent $x \in M$, let i be the first place of X in the subsequence $X^r + X$. Then M_i has the form $X + M + X$ as X^r unravels to X and Ehrenfeucht value $x + m + x$, which is x by Theorem 2.2. Thus the level t Ehrenfeucht values are precisely the persistent $x \in M$. □

Sh:432

2.5. The Ehrenfeucht Game

The Ehrenfeucht Game is a very general method for showing that two models have the same first order properties up to quantifier depth t. We concentrate on a specific example. Consider the space ΣA of finite sequences of elements of A and two models $M = \langle [u], \leq, f \rangle$, $M' = \langle [u'], \leq, f' \rangle$. The Ehrenfeucht Game has two players, Spoiler and Duplicator. There are t rounds. On each round Spoiler first selects one term from either model (i.e., either an $x \in [u]$ or an $x' \in [u']$) and then Duplicator chooses a term from the other model. Let $i_1, \ldots, i_t$ be the terms chosen from M and $i'_1, \ldots, i'_t$ be the terms chosen from M', both in the order of the rounds chosen. For Duplicator to win he must first assure that $f(i_q) = f'(i'_q)$. (Thinking of M, M' as sequences if Spoiler picks an α, then Duplicator must also pick an α.) Also he must assure that $i_q = i_r$ if and only if $i'_q = i'_r$. (When Spoiler picks a term already picked, Duplicator must pick its counterpart.) Finally he must assure that $i_q < i_r$ if and only if $i'_q < i'_r$. (So, e.g., if Spoiler picks i_7 between i_2 and i_5, then Duplicator must pick an i'_7 between i'_2 and i'_5.) If he does all this, Duplicator wins; otherwise Spoiler wins. This is a finite perfect information game with no draws. As such, someone is the winner.

Property Ehrenfeucht. *Duplicator wins the t round Ehrenfeucht Game on M, M' if and only if M, M' satisfy precisely the same sentences of quantifier depth at most t in the linear language for A.*

Now we give a reduction theorem. Consider two models $M = \langle [a], \leq, U \rangle$ and $N = \langle [b], \leq, V \rangle$ of a unary predicate. (Alternately, a sequence of zeroes and ones.) Consider decompositions $M = M_1 + \cdots + M_u$ and $N = N_1 + \cdots + N_v$. Let $m_1, \ldots, m_u$ and $n_1, \ldots, n_v$ be the t-level Ehrenfeucht values of the respective M's and N's. Suppose all the m_i, $n_j \in P$, a finite set. Suppose $m_1 \cdots m_u$ and $n \cdots n_v$ are equivalent as elements of ΣP, in the sense that they satisfy the same sentences up to quantifier depth t.

Theorem 2.11. *Under the above assumptions M, N satisfy the same sentences up to quantifier depth t.*

Proof (Outline). Consider the t move Ehrenfeucht game on M, N. While the game is progressing, Duplicator imagines a "supergame" on the sequences $m_1 \cdots m_u$ and $n_1 \cdots n_v$. A selection of $x \in M_i$ corresponds to selecting m_i and of $y \in N_j$ to selecting n_j. Suppose Spoiler picks $x \in M_i$, selecting in N_j being similar. Duplicator calculates that in the supergame a winning response to m_i is n_j. Duplicator will select a $y \in N_j$. But which one? The first time $x \in M_i$ is played creates a link between M_i and N_j. Duplicator imagines a "subgame" on M_i, N_j of t moves, a game that he wins since $m_i = n_j$. Whenever Spoiler plays in M_i or N_j, Duplicator plays in the other according to the subgame strategy. □

Now suppose $M = \langle [a], C, U \rangle$, $N = \langle [b], C, V \rangle$ are circular models. Suppose M can be decomposed into intervals $M_1, \ldots, M_r$ (so that the first element of M_1 immediately follows the last element of M_r) and N can be similarly decomposed into intervals $N_1, \ldots, N_r$. Let m_i, n_j be the t-level Ehrenfeucht values of the respective M_i, N_j, all lying in P. Suppose $m_1 \cdots m_u$ and $n_1 \cdots n_v$ are equivalent as elements of $Cyc(P)$, in the sense that they satisfy the same sentences up to quantifier depth t.

Theorem 2.12. *Under the above assumptions M, N satisfy the same sentences up to quantifier depth t in the circular language for A.*

The strategy for Duplicator is basically the same.

2.6. Markov Chains

Fix, for each $a \in A$, a value $p_a \in (0, 1)$ such that $\sum_{a \in A} p_a = 1$. By a random sequence on length l we mean a sequence $a_1 \cdots a_l$ with each a_i chosen independently and $\Pr[a_i = a] = p_a$ for all i and all $a \in A$. Let $M[\alpha, l]$ denote the probability that $a_1 \cdots a_l$ have Ehrenfeucht value l.

On M we define a Markov Chain by setting the transition probability from m to $m + a$ to be p_a for all $m \in M$, $a \in A$. Consider the Markov Chain to start at O at time zero. Then $M[\alpha, l]$ is precisely the probability that the Markov Chain is in state α at time l.

Any $y \in M$ represents some finite string $a_1 \cdots a_r \in \Sigma A$ so that $y = a_1 + \cdots + a_r$. For any $x, y \in M$ there is therefore a path from x to y in the Markov Chain. If $x \in M$ is persistent, then $R_x = \{x + y\colon y \in M\}$ is a strongly connected component of the Markov Chain. If x is transient, there is a path from x to some y from which it is impossible to return to x. Hence: *$x \in M$ is persistent precisely when it is a persistent state in the Markov Chain.* Property C implies that no strongly connected component can be periodic. Hence M is a finite aperiodic Markov Chain. We use a basic result from Markov Chain Theory.

Property E. *If $\alpha \in M$ is persistent, there exists $c_\alpha > 0$ so that*

$$\lim_{l \to \infty} M[\alpha, l] = c_\alpha \,.$$

If $\alpha \in M$ is transient there exist $K < 1$ and c so that for all l

$$M[\alpha, l] < cK^l \,.$$

Our application requires a more powerful result that these probabilities are not altered by small perturbations. We add a parameter p and let $p_a(p) \in (0, 1)$ depend on p, still holding $\sum_{a \in M} p_a(p) = 1$ for each p. Further assume that, for each $a \in A$, $\lim_{p \to 0} p_a(p) = p_a$. Let M_p denote the Markov Chain with transition probabilities $p_a(p)$ and let M^0 be the "limit" Markov Chain with transition probabilities p_a. Let $M_p[\alpha, l]$ and $M^0[\alpha, l]$ denote the probabilities that the respective chains, beginning at O at time zero, are at state α at time l.

Property E$^+$. *For each fixed l*

$$\lim_{p \to 0} M_p[\alpha, l] = M^0[\alpha, l] \,.$$

Sh:432

If $\alpha \in M$ is persistent,

$$\lim_{p\to 0}\lim_{l\to\infty} M_p[\alpha, l] = \lim_{l\to\infty} M^0[\alpha, l],$$

which is defined and positive by Property E. If $\alpha \in M$ is transient, there exists $p_0 > 0$, $K < 1$, and c so that, for all $p < p_0$, l

$$M_p[\alpha, l] < cK^l .$$

3. AN INFINITE MODEL

Here we consider a random unary predicate $U = U_p$ defined on the set $N = \{1, 2, \ldots\}$ of all positive integers and with $\Pr[U(i)] = p$ for all $i \in N$, the events $U(i)$ being mutually independent. Our definitions will apply for any $p \in (0, 1)$, but we note that our analysis will center on the asymptotics as $p \to 0$. All definitions and results will be relative to a fixed integer t. The definitions of the next section are, formally, independent of p.

3.1. k-Intervals

Definition. *The* 1-interval *of i is $[i, i_1)$ where $i_i > i$ is the least integer with $U(i_1 - 1)$.*

In dynamic language, and considering U as a sequence of zeroes and ones, to find the 1-interval one starts at i and keeps going to the "right" until finding a one. Now set $s = 3^t$. All 1-intervals consist of a string of zeroes (possibly empty) followed by a one.

Definition. *The* 1-value *of i is the symbol a_j when the* 1*-interval of i consists of j zeroes followed by a one and $j < s$. The* 1*-value is the symbol b if the* 1*-interval consists of $j \geq s$ zeroes followed by a one.*

From Property C, the 1-value of i determines the Ehrenfeucht value of the 1-interval of i. That is, if i, i' have the same 1-value, then their 1-intervals have the same first order properties up to quantifier depth t.

Definition. *We define the k-*interval *of i, the k-*value *of i, T_k and P_k. The definitions are done by induction on k. The case $k = 1$ has already been done; assume inductively that they have been given for k. Beginning at $i = i_0$, let $[i_0, i_1)$ be the k-interval of i_0 and then take successive k-intervals $[i_1, i_2)$, $[i_2, i_3), \ldots, [i_{u-1}, i_u)$ until reaching a k-interval $[i_u, i_{u+1})$ whose k-value lies in T_k. (This could happen at $u = 0$.) The $(k+1)$-interval of i is then $[i, i_{u+1})$. Let $x_1, \ldots, x_u, y$ be the k-values of the successive intervals so that $x_i \in P_k$ and $y \in T_k$. Now consider $x_1 \cdots x_u$ as a string, an element of ΣP_k, and let α denote its Ehrenfeucht value, as defined in Section 2.1. The $k+1$-value of i is then the pair (α, y). Let P_{k+1} denote the set of pairs (α, y) with $y \in T_k$ and α a* persistent *Ehrenfeucht class of ΣP_k. Let T_{k+1} denote the set of such pairs where α is a transient Ehrenfeucht class as defined in Section 2.3.*

Sh:432

Definition. *A k-interval $[i, j)$ is called k-persistent (or, simply, persistent) if it has k-value in P_k; otherwise it is called k-transient, or, simply, transient.*

Remarks. As this definition is somewhat the key to our entire program, several comments are in order. The rough idea is to capture events that occur every $\Theta(p^{-k})$—k consecutive ones being a natural, but by no means the only example. A persistent k-interval starting at i ends when a "typical" event of probability $\sim p^k$ occurs; the transient k-intervals are when something "atypical" occurs.

Example. Take $t = 5$, $s = 243$. A persistent 1-interval consists of at least 243 zeroes followed by a one. $P_1 = \{b\}$. In ΣP_1 a sequence is persistent if it consists of more than 243 b's. So a persistent 2-interval looks like at least 243 persistent 1-intervals followed by $a_i - i - 1$ zeroes and then another one. A persistent 2-interval ends with two ones close together. The persistent states can be denoted B_i, $1 \le i \le 243$. A 2-interval with value B_i consists of "many" ones "far" apart followed by two ones i apart. A typical transient 2-value, let us denote it by $C_{8,3}$, consists of eight persistent 1-intervals, two zeroes, and a one. The "atypical" thing is that the two ones close together come too soon. The general transient 2-value is of the form $C_{a,b}$ with $0 \le a < 243$ and $0 \le b < 243$.

The 3-intervals introduce the real complexities. How can a persistent 3-interval end? It ends with a persistent 2-interval, which ends with two ones close together, followed by a transient 2-interval. If the latter is of the form $C_{0,b}$, then there are three ones close together. If the latter is of the form $C_{a,b}$, then there are two pairs of close together ones and between them only a ones all far apart. When is a 3-interval persistent? Suppose its persistent 2-intervals have 2-values $B_{i_1} \cdots B_{i_u}$. We must have that string persistent in ΣP_2. As an example, the 3-interval would be transient if no B_{23} appeared in the string—i.e., there were not two ones precisely 23 apart. But similarly for it to be persistent (for, say, $t \ge 3$) there must be two B_{23} with no B_{86} between them.

The open interval, $[i, \infty)$ is split, for every k, into an infinite sequence of k-intervals $I_1^k, I_2^k, \ldots$. Each I_l^k is the union of consecutive $(k-1)$-intervals. Many things can "cause" a k-interval to end. Here we give the most natural: a sequence of k ones.

Lemma. *If $U(x)$, then for every $k \ge 1$ the k-interval of x is $[x, x+1)$ and is transient.*

Proof. Induction on k. The case $k = 1$ follows directly from the definitions. Assume for $k - 1$. As $[x, x+1)$ is a transient $(k-1)$-interval, $[x, x+1)$ is the k-interval of x. Its k-value is (O, y) with y the $(k-1)$-value of $[x, x+1)$ and O the Ehrenfeucht value of the null sequence in ΣP_{k-1}. But O is certainly transient so $[x, x+1)$ is a transient k-interval. □

Theorem 3.1. *Let $i \le j$ and suppose $U(s)$ for $j \le s \le j + k - 1$. Decompose $[i, \infty)$ into consecutive k-intervals $I_1^k, I_2^k, \ldots$. Then one of the k-intervals has final value $j + k - 1$.*

Proof. Induction on k. For $k = 1$ this is immediate. Assume for $k - 1$ so that when $[i, \infty)$ is decomposed into $(k-1)$-intervals some interval $I_l^{(k-1)}$ ends in

Sh:432

$j+k-2$. By the Lemma, $I_{l+1}^{(k-1)} = [j+k-1, j+k)$ and is transient. Hence $I_{l+1}^{(k-1)}$ ends a k-interval. □

3.2. Probability

Now we examine the probabilities of the various k-values for the random unary predicate. All asymptotics are as $p \to 0$.

Definition. *For any k-value β, let $P[\beta]$ denote the probability that the k-interval of i has k-value β.*

Note that $P[\beta]$ is independent of i. Formally we should write $P[\beta, p]$, but we suppress the p in this and later functions for notational convenience.

Theorem 3.2. *If β is persistent, there exists $c_\beta > 0$ with*

$$P[\beta] = c_\beta + o(1) .$$

If β is transient, then there exists $c_\beta > 0$ with

$$P[\beta] = c_\beta p + o(p) .$$

Proof. The proof is by induction on k. For $k = 1$ and $\beta = a_j$ (i.e., $j-1$ zeroes followed by a one) $P[\beta] = (1-p)^{j-1}p \sim p$. For $\beta = b$ (at least s zeroes followed by a one)

$$P[\beta] = 1 - \sum_{j=1}^{s} (1-p)^{s-1}p = 1 - o(1) .$$

Now assume the result for $k-1$. Let γ be the probability that the $(k-1)$-value of i is transient. By induction $\gamma = cp + o(p)$ since γ is the (finite) sum of the probabilities of the $(k-1)$-value being β over all transient β. For all persistent $(k-1)$-values β set

$$P^P[\beta] = P[\beta]/(1-\gamma) ,$$

the conditional probability of a $(k-1)$-value being β give that it is persistent. As $\gamma = o(1)$, $P^P[\beta] = c_\beta + o(1)$. For all transient y set

$$P^T[y] = P[y]/\gamma ,$$

the conditional probability of a $(k-1)$-value being y given that it is transient. As $\gamma = cp + o(p)$, $P^T[y] = c_y^T + o(1)$ with $c_y^T = c_y/c$.

Now comes an essential point. Let the successive $(k-1)$-intervals for i have values $\beta_1, \beta_2, \ldots$. Conditioning on $\beta_1, \ldots, \beta_u$—even conditioning on the precise sequence giving these values—the $(k-1)$-value for the next interval is still independent. That is, having examined the sequence up to a certain point gives us

Sh:432

no change in the distribution of the sequence after that point. Hence β_{u+1} is independent of $\beta_1, \ldots, \beta_u$.

Let $\beta = (\alpha, y)$ be a k-value. For $l \geq 0$ consider a random string $\beta_1 \cdots \beta_l$ of elements of P_{k-1}, independently chosen each with distribution P^P. Let $M[\alpha, l]$ denote the probability that this random string, as an element of ΣP_{k-1} has Ehrenfeucht value α. We claim:

$$P[\beta] = \sum_{l=0}^{\infty} (1-\gamma)^l M[\alpha, l] \gamma P^T[y]$$

For the k-interval of i to have value β the successive $(k-1)$-intervals must have, for some $l \geq 0$, persistent values $\beta_1 \cdots \beta_l$ followed by transient y. Conditioning on persistency and transience gives the $(1-\gamma)^l$ and γ factors, respectively. Under that conditioning, $M[\alpha, l]$ and $P^T[y]$ are the probabilities of getting α and y, respectively.

Assume $\beta = (\alpha, y)$ is persistent. Set $c_y = \lim_{p\to 0} P^T(y)$ and, by Property E^+, $c_\alpha = \lim_{p\to 0} \lim_{l\to\infty} M[\alpha, l]$. As $\gamma \to 0$ and M, P^T are uniformly bounded, standard analysis gives

$$\lim_{p\to 0} P[\beta] = c_y c_\alpha .$$

Now assume $\beta = (\alpha, y)$ is transient. By Property E^+, $M[\alpha, l] < cK^l$ for all $p < p_0$. Thus for any $\epsilon > 0$ there exists l_0 such that

$$\sum_{l=l_0}^{\infty} (1-\gamma)^l M[\alpha, l] P^T[y] < \sum_{l=l_0}^{\infty} cK^l < \epsilon$$

uniformly for all $p < p_0$. Recall M is here also a function of p. Set $M^0[\alpha, l] = \lim_{p\to 0} M[\alpha, l]$, using property E^+. Then

$$\lim_{p\to 0} P[\beta]\gamma^{-1} = \lim_{l_0\to\infty} \lim_{p\to 0} \sum_{l<l_0} (1-\gamma)^l M[\alpha, l] P^T[y]$$

$$= \lim_{l_0\to\infty} \sum_{l<l_0} M^0[\alpha, l] c_y = c_y \sum_{l=0}^{\infty} M^0[\alpha, l] .$$

As $\gamma = cp + o(p)$,

$$P[\beta] = \left[cc_y \sum_{l=0}^{\infty} M^0[\alpha, l] \right] p + o(p)$$

as desired. □

Remark. Note that the contribution to $P[\beta]$ is dominated by l large when β is persistent and by l small when β is transient.

Example. Take $t = 5$, $s = 243$, $A = \{a, b, c\}$, $k = 2$. The persistent 2-intervals B_i consist of "many" ones "far" apart and then two ones i apart, $1 \leq i \leq 243$. Each one is asymptotically equally likely and $P[B_i] \to 1/243$.

3.3. Lengths of k-Intervals

Let $X_k = X_k(p, i)$ denote the length of the k-interval beginning at i. Clearly the distribution of X_k does not depend on i. Here we make precise the notion that k-intervals measure events occurring every $\sim p^k$.

Theorem 3.3. *For each $k \geq 1$, there exist positive p_k, ϵ_k, ϵ_k' so that for $p < p_k$*

$$\Pr[X_k > \epsilon_k p^{-k}] > \epsilon_k'$$

Proof. First set $k = 1$. A 1-interval ends with a one so that $\Pr[X_1 \leq a] \leq pa$ and we can set $p_1 = \epsilon_1 = \epsilon_1' = 0.1$ with room to spare. Now we use induction, assuming the result for k. Let γ be the probability that a k-interval is transient. From the proof of Theorem 3.2, $\gamma(p) = c_1 p + o(p)$ for a positive c_1. Pick any $c > c_1$ and so that for p sufficiently small $\gamma < cp$.

Let X_{k+1}' be the sum of the lengths of the first $p^{-1}/2c$ k-intervals, beginning at i. Each each has independent length so almost surely (as $p \rightarrow 0$) $(\epsilon_k' + o(1))p^{-1}/2c$ of them have length at least $\epsilon_k p^{-k}$ and $X_k' > (\epsilon_k \epsilon_k'/2c + o(1))p^{-(k+1)}$. For p sufficiently small $X_k' > (\epsilon_k \epsilon_k'/4c)p^{-(k+1)}$ with probability at least 0.9. The probability that any of the first $p^{-1}/2c$ k-intervals is transient is less than $\gamma p^{-1}/2c < 0.6$. At least 30% of the time the first occurs and the second does not. But then $X_k > X_k'$. Thus we may take $\epsilon_{k+1}' = 0.3$ and $\epsilon_{k+1} = \epsilon_k \epsilon_k'/4c$. □

Definition. *Consider the sequence $I_1^k, I_2^k, \ldots$ of k-intervals beginning at $i = 1$. Let $L = L_k(n, p)$ denote the number of these intervals lying entirely in $[1, n]$.*

Theorems 3.4. *Let $n \rightarrow \infty$, $p \rightarrow 0$ so that $np^k \rightarrow \infty$ and $np^{k+1} \rightarrow 0$. Then*

(i) *$L_k(n, p) > np^k(1 - o(1))$ almost surely.*
(ii) *$L_{k+1}(n, p) \rightarrow 0$ almost surely.*

Remark. Indeed, even more is true. $L_k(n, p) = \Theta(np^k)$ and even the proper constant may be determined.

Proof. Basic probability arguments given that the random sequence in $[1, n]$ almost surely contains k-consecutive ones $\sim np^k$-times. By Theorem 3.1 each such sequence ends a k-interval. For (ii), let ϵ_k, ϵ_k' be given by Theorem 3.3. Consider the first $2np^k/\epsilon_k \epsilon_k'$ k-intervals, beginning at one. As $np^k \rightarrow \infty$ almost surely $\sim 2np^k/\epsilon_k$ of them have length at least $\epsilon_k p^{-k}$ so their total length is more than, asymptotically, $2n$ and hence almost surely more than n. The probability that any of them is transient is less than $\gamma(2np^k/\epsilon_k \epsilon_k')$. As $\gamma = \Theta(p)$, this approaches zero. Hence almost surely the first $2np^k/\epsilon_k \epsilon_k'$ k-intervals are all transient and reach past n so that the $(k+1)$-interval beginning at one is not ended by n. □

3.4. The Glue

With $np^{k+1} \rightarrow 0$ and $np^k \rightarrow \infty$, we now have a picture of $U_{n,p}$ as a long succession of k-intervals $I_1^k, I_2^k, \ldots, I_Q^k$ followed by some excess—the "cutoff" part of I_{Q+1}^k.

In this somewhat technical section we show how to glue this excess onto the front of I_1^k.

Definition. *A persistent k-interval I beginning at one is called* superpersistent *if the following holds for* $1 \le w < k$: *Let* $I^1, I^2, \ldots, I^L, I^{L+1}$ *be the successive w-intervals beginning at one with* I^{L+1} *the first transient one and let* $\beta_1, \ldots \beta_L, y$ *denote their respective Ehrenfeucht values. Then we require that* $\beta_2 \cdots \beta_L$ *be persistent in* ΣP_w. *(When* $k = 1$, *all persistent intervals are called superpersistent.)*

Remark. As the full k-interval is persistent, the first $(w+1)$-interval is persistent and so $\beta_1 \cdots \beta_L$ is persistent. Superpersistency only further requires persistency of the string without β_1.

Theorem 3.5. *Let J be a model on* $[-u, 0]$ *so that the k-interval beginning at* $-u$ *does not end by* 0. *Let I be a superpersistent k-interval on* $[1, a]$. *Then* $J + I$ *is a persistent k-interval on* $[-u, a]$.

Proof. When $k = 1$, J is a sequence of all zeroes and I has "many" zeroes followed by a one so that $J + I$ also has "many" zeroes followed by a one, and is persistent. Now we use induction. Split J into $(k-1)$-intervals, number them J^{-b}, $J^{-b+1}, \ldots, J^0$ with excess J^*. [If a $(k-1)$-interval was just completed, then J^* will be null.] Let $\beta_{-b}, \ldots, \beta_0$ denote their $(k-1)$-values. Split I into $(k-1)$-intervals $I^1, \ldots, I^{L+1}$ with $(k-1)$-values $\beta_1, \ldots, \beta_L, y$. Observe that I being a superpersistent k-interval implies that I^1 is a superpersistent $(k-1)$-interval. Observe that J^*, as excess, does not complete a $(k-1)$-interval. By induction $J^* + I^1$ is a persistent $(k-1)$-interval. Its $(k-1)$-value, call it β_1', may be different from β_1. The $(k-1)$-intervals of $J + I$ thus have values $\beta_{-b}, \ldots, \beta_0$, $\beta_1'\beta_2, \ldots, \beta_L, y$. All are persistent but y so $J + I$ is a k-interval. By superpersistency $\beta_2 \cdots \beta_L$ is persistent in ΣP_{k-1}. By Theorem 2.3 the addition of *any* prefix retains persistency so $\beta_{-b} \cdots \beta_0, \beta_1'\beta_2 \cdots \beta_L$ is persistent in ΣP_{k-1} and hence $J + I$ is a persistent k-interval. □

Theorem 3.6. *Almost surely, as* $p \to 0$, *the k-interval beginning at one is superpersistent.*

Proof. Fix $1 \le w < k$. Let $\beta_1, \beta_2, \ldots, \beta_L, y$ denote the Ehrenfeucht values of the successive i-intervals beginning at one, with y the first transient one. From Theorem 3.2, the first $(w+1)$-interval is almost surely persistent so that $\beta_1 \cdots \beta_L$ is almost surely persistent in ΣP_w. Also from Theorem 3.2, β_1 is almost surely not transient. Conditioning on β_1 being not transient, the distribution of $\beta_2 \cdots \beta_L$ (i.e., stopping at the first transient) is the same as the distribution of $\beta_1 \cdots \beta_L$ without the conditioning and so almost surely $\beta_2 \cdots \beta_L$ is persistent. An event that occurs almost surely when conditioned on an almost sure event occurs almost surely. As this holds for each of the finite number of i, the k-interval beginning at one is almost surely superpersistent.

4. THE ZERO–ONE LAW FOR CIRCULAR SEQUENCES

Now we can give the proof of Theorem 1.1. Fix k. Let A be any first order sentence. Let t be the quantifier depth of A. Fix $p = p(n)$ with $np^k \to \infty$ and $np^{k+1} \to 0$. Now consider the random unary predicate $U_{n,p}$. Split U into consecutive k-intervals $I_1^k, \ldots, I_Q^k$ followed by an excess J. (That is, take the infinite model and let J be I_{Q+1}^k cut off at n.) Almost surely I_1^k is superpersistent. Placing, by Theorem 3.5, J in front of I_1^k we decompose U into k-intervals $I_1^{k*}, I_2^k, \ldots, I_Q^*$. By Theorem 3.4 almost surely $Q > (1 - o(1))np^k$. Set $R = \lfloor (np^k)^{1/2} \rfloor$ so that $R \to \infty$ and $R < Q$ almost surely. Let $\beta_2, \ldots, \beta_R$ be the k-values of $I_2^k, \ldots, I_R^k$. These are independently distributed; each value $\beta \in P_k$ occurs with asymptotically positive distribution, so as $R \to \infty$ the sequence $\beta_2 \cdots \beta_R$ almost surely contains the universal sequence σ (whose size is *fixed*, given k, t) of Theorem 2.10. But once σ is a consecutive subsequence Theorems 2.10 and 2.12 assure that the truth value of all first order sentences of quantifier depth at most t, in particular our desired A, are determined.

5. CONVERGENCE FOR THE LINEAR MODEL

We have already remarked in Section 1 that Zero–One Laws generally do not hold for the linear model $\langle [n], \leq, U \rangle$ and that Dolan has characterized those $p = p(n)$ for which they do. Our main object in this section is the following *convergence* result.

Theorem 5.1. *Let k be a positive integer, and S a first order sentence. Then there is a constant $c = c_{k,S}$ so that, for any $p = p(n)$ satisfying*

$$np^k \to \infty \quad \text{and} \quad np^{k+1} \to 0\,,$$

we have

$$\lim_{n \to \infty} \Pr[U_{n,p} \models S] = c\,.$$

Again we shall fix the quantifier depth t of S and consider Ehrenfeucht classes with respect to that t. For each $\beta \in P_k$, let c_β, as given by Theorem 3.2, be the limiting probability that a k-interval has k-value β. Let M be the set of equivalence classes of ΣP_k, a Markov Chain as defined in Section 2.4, and for each $\equiv_R$-class R_x let $P[R_x]$ be the probability that a random sequence $\beta_1 \beta_2 \cdots$ eventually falls into R_x.

In $\langle [n], \leq, U \rangle$ let $\beta_1 \cdots \beta_N$ denote the sequence of k-values of the successive k-intervals, denoted $[1, i_1), [i_1, i_2), \ldots$, from 1.

Set, with foresight, $\delta = 10^{-2} 3^{-t}$.

We shall call U on $[n]$ *right nice* if it satisfies two conditions. The first is simply that all the $\beta_1, \ldots, \beta_N$ described above are persistent. Theorem 3.4 gives that this holds almost surely. The second will be a particular universality condition. Let $A_1 \cdots A_R$ be a specific sequence in ΣP_k with the property that for every R_x and

Sh:432

L_y there exists a q so that

$$A_1 \cdots A_1 \in L_y \quad \text{and} \quad A_{q+1} \cdots A_R \in R_x .$$

(We can find such a sequence for a particular choice of R_x and L_y by taking specific sequences in ΣP_k in those classes and concatenating them. The full sequence is achieved by concatenating these sequences for all choices of R_x and L_y. Note that as some $A_1 \cdots A_q \in L_y$ the full sequence is persistent.) The second condition is that inside any interval $[x, x+\delta n] \subset [1, n]$ there exist R consecutive k-intervals $[i_L, i_{L+1}), \ldots, [i_{L+R}, i_{L+R+1})$ whose k-values are, in order, precisely $A_1, \ldots, A_R$. We claim this condition holds almost surely. We can cover $[1, n]$ with a finite number of intervals $[y, y+(\delta/3)n]$ and it suffices to show that almost always all of them contain such a sequence, so it suffices to show that a fixed $[y, y+(\delta/3)n]$ has such a sequence. Generating the k-intervals from 1 almost surely a k-interval ends after y and before $y+(\delta/6)n$. Now we generate a random sequence $\beta_1 \cdots$ on an interval of length $(\delta/6)n$. But constants do not affect the analysis and almost surely $A_1 \cdots A_R$ appears.

Now on $\langle [n], \leq, U \rangle$ define U^r by $U^r(i)$ if and only if $U(n+1-r)$. U^r is the sequence U in reverse order. Call U left nice if U^r is right nice. Call U nice if it is right nice and left nice. As all four conditions hold almost surely, the random $U_{n,p}$ is almost surely nice.

Let U be nice and let $\beta_1 \cdots \beta_N$ and $\beta_1^r \cdots \beta_{N^r}^r$ denote the sequences of k-values for U and U^r respectively and let R_x and R_{x^r} denote their $\equiv_R$-classes, respectively. (Both exist since the sequences are persistent.)

Theorem 5.2. *The values R_x and R_{x^r} determine the Ehrenfeucht value of nice U.*

Proof. Fix two models $M = \langle [n], \leq, U \rangle$ and $M' = \langle [n'], \leq, U' \rangle$, both nice and both with the same values R_x, R_{x^r}. Consider the t-move Ehrenfeucht game. For the first move suppose Spoiler picks $m \in M$. By symmetry suppose $m \leq n/2$. Let $[i_{r-1}, i_r)$ be one of the k-intervals with, say, $0.51n \leq i_r \leq 0.52n$. We allow Duplicator a "free" move and have him select i_r. Let $\beta_1 \cdots \beta_N$ and $\beta'_1 \cdots \beta'_{N'}$ be the sequences of k-values for M and M', respectively. Let z be the class of $\beta'_1 \cdots \beta_r$. Since U is nice, this sequence already contains $A_1 \cdots A_R$ and hence is persistent so $z \in R_x$. Let z' be the class of $\beta_{r+1} \cdots \beta_N$. By the same argument z' is persistent. In M' inside of, say, $[0.5n, 0.51n]$ we find the block $A_1 \cdots A_R$. By the universality property we can split this block into a segment in L_z and another in $R_{z'}$. Adding more to the left or right doesn't change the nature of this split. Thus there is an interval $[i'_{r'-1}, i'_{r'})$ so that $\beta'_1 \cdots \beta'_{r'} \in L_z$ and $\beta'_{r'+1} \cdots \beta'_{N'} \in R_{z'}$. Spoiler plays $i'_{r'}$ in response to i_r.

The class of $\beta_1 \cdots \beta_r$ is z and $z \in R_x$. The class z' of $\beta'_1 \cdots \beta'_{r'}$ is in L_z and R_x. As $z \in L_z \cap R_x$, $z = z'$. Thus $[1, i_r)$ under M and $[1, i'_{r'})$ under M' have the same Ehrenfeucht value. Thus Duplicator can respond successfully to the at most t moves (including the initial move m) made in these intervals. Thus Spoiler may as well play the remaining $t-1$ moves on $M_1 = \langle [i_r, n], \leq, U \rangle$ and $M'_1 = \langle [i'_{r'}, n'], \leq, U' \rangle$. These intervals have lengths $n_1 \geq n/3$ and $n'_1 \geq n'/3$, respectively. But now M and M' are both nice with respect to $\delta_1 = 3\delta$—the sequence $A_1 \cdots A_R$ still appears inside every interval of length $\delta n \leq \delta_1 n_1$ in M and $\delta_1 n'_1$ in

Sh:432

M'. Hence we can apply the same argument for the second move—for convenience still looking at Ehrenfeucht values with respect to the t move game. After t moves we still have nice M_t, M'_t with respect to $\delta_t \le 10^{-2}$ so the arguments are still valid. But at the end of t rounds Duplicator has won. □

Proof of Theorem 5.1. Let R_x, $R_{x'}$ be any two $\equiv_R$-classes. Let U be random and consider $\langle [\delta n\rangle, \le, U\rangle$. The sequence of k-values lies in R_x with probability $P[R_x] + o(1)$. The same holds for U^r on $[\delta n]$. But U^r examines U on $[(1-\delta)n, n]$ so as $\delta < 0.5$ the values of the $\equiv_R$-classes are *independent* and so the joint probability of the values being R_x and $R_{x'}$, respectively, is $P[R_x]P[R_{x'}] + o(1)$. Given Theorem 5.2, $\langle [n], \le, U\rangle$ then has a value $v = v(R_x, R_{x'})$. As

$$\sum P[R_x]P[R_{x'}] = \sum P[R_x] \sum P[R_{x'}] = 1 \times 1 = 1\,,$$

this gives a limiting distribution for the Ehrenfeucht value v on $\langle [n], \le, U\rangle$. □

REFERENCES

[1] P. Dolan, A. zero-one law for a random subset, *Random Struct. Alg.* **2**, 317–326 (1991).

[2] A. Ehrenfeucht, An application of games to the completeness problem for formalized theories, *Fundam. Math.*, **49**, 129–141 (1961).

[3] T. Łuczak and J. Spencer, When does the zero-one law hold? *J. Am. Math. Soc.*, **4**, 451–468 (1991).

[4] S. Shelah and J. Spencer, Zero-one laws for sparse random graphs, *J. Am. Math. Soc.*, **1**, 97–115 (1988).

Received April 24, 1992
Revised December 15, 1992
Accepted July 7, 1993