Logic and Random Structures

Joel Spencer

> In the world of randomization almost everything seems to be possible. – Michael Rabin

1 An Instructive Example

We begin with a rather easy random model which illustrates many of the concepts we shall deal with. We call it the simple unary predicate with parameters n, p and denote it by $SU(n,p)$. The model is over a universe Ω of size n, a positive integer. We imagine each $x \in \Omega$ flipping a coin to decide if $U(x)$ holds, and the coin comes up heads with probability p. Here we have p real, $0 \leq p \leq 1$. Formally we have a probability space on the possible U over Ω defined by the properties $\Pr[U(x)] = p$ for all $x \in \Omega$ and the events $U(x)$ being mutually independent. We consider sentences in the first order language. In this language we have only equality (we shall always assume we have equality) and the unary predicate U. (The cognescenti should note that Ω has no further structure and in particular is not considered an ordered set.)

This is a rather spartan language. One thing we can say is

$$YES := \exists_x U(x),$$

that U holds for some $x \in \Omega$. Simple probability gives

$$\Pr[SU(n,p) \models YES] = 1 - (1-p)^n$$

As p moves from zero to one $\Pr[YES]$ moves monotonically from zero to one. We are interested in the asymptotics as $n \to \infty$. At first blush this seems trivial: for $p = 0$, $SU(n,p)$ never models YES while for any constant $p > 0$,

$$\lim_{n\to\infty} \Pr[SU(n,p) \models YES] = \lim_{n\to\infty} 1 - (1-p)^n = 1$$

In an asymptotic sense YES has already almost surely occured by the time p reaches any positive constant.

This leads us to a critical notion. *We do not restrict ourselves to p constant but rather consider $p = p(n)$ as a function of n.* What is the parametrization $p = p(n)$ that best enables us to see the transformation of $\Pr[SU(n,p(n)) \models YES]$ from zero to one. Some reflection leads to the parametrization $p(n) = c/n$. If c is a positive constant then

$$\lim_{n\to\infty} \Pr[SU(n,p(n)) \models YES] = \lim_{n\to\infty} 1 - (1 - \frac{c}{n})^n = 1 - e^{-c}$$

(Technically, as $p \leq 1$ always, this parametrization is not allowable for $n < c$ - but since we are only concerned with limits as $n \to \infty$ this will not concern us.) If we think of c going from zero to infinity then the limit probability is going from zero to one. We are actually less interested (in this exposition) in the actual limits than in whether the limits are zero or one.

We say that a property A holds *almost always* (with respect to a given $p(n)$ if $\lim_{n\to\infty} \Pr[SU(n,p(n)) \models A] = 1$. We say that A holds *almost never* if the above limit is zero or, equivalently, if $\neg A$ holds almost surely. This notion is extremely general. Whenever we have for all sufficiently large positive integers n a probability space over models of size n then we can speak of a property A holding almost surely or almost never. For the particular property YES the exact results above have the following simple consequences:

- If $p(n) \ll n^{-1}$ then YES holds almost never.
- If $p(n) \gg n^{-1}$ then YES holds almost surely.

Thus, for example, when $p(n) = n^{-1.01}$ YES holds amost never while when $p(n) = n^{-0.99}$ YES holds almost surely.

We shall say n^{-1} is a *threshold function* for the property YES. More generally, suppose we have a notion of a random model on n vertices with probability p of some predicate. We say $p_0(n)$ is a threshold function for a property A if whenever $p(n) \ll p_0(n)$ the property A holds almost never and whenever $p(n) \gg p_0(n)$ then the property A holds almost surely. This notion, due to Paul Erdős and Alfred Rényi, says roughly that $p_0(n)$ is the "region" around which $\Pr[A]$ is moving from near zero to near one. The threshold function, when it exists, is not totally determined - we could have taken $5/n$ as the threshold function for YES - but is basically determined up to constant factors. In a rough way we think of $p(n)$ increasing through the functions of n - e. g. from n^{-2} to n^{-1} to $n^{-1}\ln n$ to $\ln^{-5} n$ - and the threshold function is that place where $\Pr[A]$ changes.

A natural problem for probabilists is to determine the threshold function, if one exists, for a given property A. For logicians the natural question would be to determine all possible threshold functions for all properties A expressible in a given language L. Unfortunately there are technical difficulties (especially with later more complex models) with threshold functions -

properties A need not be monotone, threshold functions need not exist, and, worst of all, the limits of probabilities might not exist. Rather, the logician looks for a *Zero-One Law* of which the following is prototypical:

Theorem: Let $p = p(n)$ satisfy $p(n) \gg n^{-1}$ and $1 - p(n) \gg n^{-1}$. Then for *any* first order property A

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

Further, the limiting value depends only on A and not on the choice of $p(n)$ within that range.

Our approach to this theorem, which shall also be used in later more complex cases, is to find an explicit theory T such that

- Every $A \in T$ holds almost surely
- T is complete

Will this suffice? When $T \models B$ finiteness of proof gives that B follows from some $A_1, \ldots, A_s \in T$ and hence from $A_1 \wedge \ldots \wedge A_s$. But the finite conjunction of events holding almost surely holds almost surely so B would hold almost surely. By completeness, either $T \models B$ or $T \models \neg B$, and in the latter case $\neg B$ holds almost surely so that B holds almost never.

In our situation T is given by two simple schema.

1. (For $r \geq 1$) There exist distinct $x_1, \ldots, x_r$ with $U(x_i)$ for $1 \leq i \leq r$.
2. (For $r \geq 1$) There exist distinct $x_1, \ldots, x_r$ with $\neg U(x_i)$ for $1 \leq i \leq r$.

Note that the number X of x with $U(x)$ has Binomial Distribution with parameters $n, p(n)$ – that the event $X \geq r$ holds almost surely follows from basic probabilistic ideas from the assumption $np(n) \to \infty$. The second schema follows from $n(1 - p(n)) \to \infty$, reversing the roles of U and $\neg U$.

Why is this T complete? Proving completeness of a theory T is bread and butter to the logic community – from the myriad of methods we choose a combinatorial approach based on the Ehrenfeucht game, as described in § 14. Let $t \geq 1$ be arbitrary and let M_1, M_2 be two countable models of T. It suffices to show that Duplicator wins the game $\mathrm{EHR}(M_1, M_2; t)$.

In our case the Duplicator strategy is simple. A countable model M of T must have an infinite number of $x \in M$ with $U(x)$ (as for all $r \geq 1$ it must have at least r such x) and, similarly, an infinite number of $x \in M$ with $\neg U(x)$. Now when Spoiler selects, say, a new $x \in M_1$ with $U(x)$ Duplicator simply selects a new $x' \in M_2$ with $U(x')$ – as there are only a finite number t of moves he cannot run out of possible x'.

In this instance the countable models of T were particularly simple - indeed the theory T was $\aleph_0$-categorical, all countable models were isomorphic. In future more complex situations this will generally not be the case and indeed we find the study of the countable models of the almost sure theory T to be quite intriguing in its own right.

2 Random Graphs

A graph G consists of a set of vertices V and an areflexive symmetric binary relation on V. We write the relation $x \sim y$ and say x, y are adjacent. Pictorially, there is an edge from x to y. For the graphtheorists, our graphs are undirected, with neither loops nor multiple edges. The random graph $G(n,p)$ ($n \geq 1$ integral, p real, $0 \leq p \leq 1$) is on a vertex set V of size n where for each distinct x, y $\Pr[x \sim y] = p$ and these events are mutually independent. We may think of each pair x, y of vertices flipping a coin to decide whether or not to have an edge between them, where the probability the coin comes up heads is p.

It is a relatively rare area of mathematics that has an explicit starting point. The subject of Random Graphs began with a monumental paper by Paul Erdős and Alfred Rényi in 1960. The very title of their paper, "On the Evolution of the Random Graph," speaks to a critical vantagepoint. As the edge probability p increases the random graph $G(n,p)$ increases in complexity. For many natural properties A there will be a threshold function $p_0(n)$ for its occurance. As in §1, when $p(n) \ll p_0(n)$ A will hold almost never while when $p(n) \gg p_0(n)$ A will hold almost always. Finding threshold functions has been a major preoccupation for researchers in Random Graphs. Lets give some examples, together with some intuitive justification for the threshold functions.

• Containing a K_4 - i. e. containing four vertices with all six pairs adjacent. The threshold function is $n^{-2/3}$. There are $\binom{n}{4} \sim n^4/24$ possible K_4s and each has the six adjacencies with probability p^6 so that the expected number of K_4s is $\sim n^4p^6/24$. When $p(n) \ll n^{-2/3}$ this expectation goes to zero so that almost surely there are none of them. When $p(n) \gg n^{-2/3}$ this expectation goes to infinity. By itself, this does not imply that almost surely there is at least one but more refined methods - in particular, an examination of the variance of the number of K_4s - do show that almost surely there will be a K_4.

• Containing a triangle. The threshold function is n^{-1} for reasons similar to those above.

• No isolated vertices. In first order language $\forall_x \exists_y x \sim y$. Here $n^{-1} \ln n$ is the threshold function. Roughly a given vertex x has probability $(1-p)^{n-1} \sim e^{-pn}$ of being isolated. When $pn > (1+\epsilon)\ln n$ this probability is $o(n^{-1})$ so that the expected number of isolated vertices is $o(1)$ and almost surely there are none. When $pn < (1-\epsilon)\ln n$ this probability is $\gg n^{-1}$ so that the expected number of isolated vertices goes to infinity and more refined techniques show that almost surely there are isolated vertices.

• Connectivity. This was one of the most beautiful results in the Erdős-Rényi paper. It turns out that connectivity has the same behavior as no isolated vertex. Their result was amazingly precise. Parametrize $p = \frac{\ln n}{n} + \frac{c}{n}$. For c any real (positive or negative) constant

$$\lim_{n \to \infty} \Pr[G(n,p) \text{ connected}] = e^{-e^{-c}}$$

• Every two vertices have a common neighbor. In first order language $\forall_{x_1} \forall_{x_2} \exists_{y_1} y_1 \sim x_1 \wedge y_1 \sim x_2$. The threshold function is $n^{-1/2} \ln^{1/2} n$. Any x_1, x_2 have an expected number $(n-2)p^2 \sim np^2$ common neighbors. This would naturally lead us to consider $p = n^{-1/2}$. Indeed, for $p \ll n^{-1/2}$ a randomly chosen x_1, x_2 will not have a common neighbor while for $p \gg n^{-1/2}$ a randomly chosen x_1, x_2 will have a common neighbor, indeed many common neighbors. But this does not suffice for *every* pair x_1, x_2 to have a common neighbor, for that one needs the extra polylogarithmic term.

• Every two vertices are joined by a path of length three. In first order language $\forall_{x_1} \forall_{x_2} \exists y_1 \exists y_2 x_1 \sim y_1 \wedge y_1 \sim y_2 \wedge y_2 \sim x_2$. The threshold function is $n^{-2/3} \ln^{1/3} n$. Any x_1, x_2 have $\binom{n-2}{2} \sim n^2/2$ potential paths (choices of y_1, y_2) of length three and each potential path has its three adjacencies with probability p^3 so that the expected number of paths is $\sim n^2 p^3/2$. This would lead us to consider $p = n^{-2/3}$ as a threshold function but, as above, an extra polylogarithmic term is needed to assure that *every* pair x_1, x_2 has such a path.

These threshold functions, and countless others, seemed to this author to have a common property: the power of n involved was always a rational number. There might be other, generally polylogarithmic, factors but they would be of smaller order than the power of n. Nowhere, so it seemed, was there a natural property with threshold function, say, $p = n^{-\pi/7}$. In 1988 this author and Saharon Shelah were able to give a formal justification for this observation and this result is the centerpiece of our discussions:

Theorem: Let $0 < \alpha < 1$, α *irrational*. Set $p(n) = n^{-\alpha}$. Then for every first order property A

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

The situation with $\alpha > 1$ has also been studied. It turns out to be considerably simpler than the $0 < \alpha < 1$ case and will not be considered here.

Our approach will be that used in §1. We shall find a theory $T = T_\alpha$ such that each $A \in T_\alpha$ holds almost surely and T_α is shown complete, using countable models and the Ehrenfeucht game. We shall need several preliminaries.

3 Extension Statements

The examples above: Every vertex has a neighbor, every two vertices have a common neighbor, every two vertices are joined by a path of length three are all examples of a vital kind of first order statements that we shall call extension statements. These statements are of the form "For all $x_1, \ldots, x_r$ there exist $y_1, \ldots, y_v$ P" where P is that certain adjacencies between some y_i, y_j and some x_i, y_j must exist. P never considers adjacencies between pairs x_i, x_j and never demands nonadjacency. We allow the case $r = 0$, so that the extension statement reduces to a purely existential statement, but require $v > 0$.

To formalize this we define a rooted graph to be a pair (R, H) where H is a graph (with $V(H), E(H)$ denoting its vertex and edge sets respectively) and R is a proper subset of the vertices. Labelling the roots $x_1, \ldots, x_r$ and the nonroots $y_1, \ldots, y_v$ we define the extension statement $Ext(R, H)$ to be that for all $x_1, \ldots, x_r$ there exist $y_1, \ldots, y_r$ having the edges of H, where we don't examine the edges between the roots and we allow extra edges. A rooted graph (R, H) has three parameters. The number of roots is denoted by r. The number of nonroots is denoted by v. The number of edges (where edges between roots are not counted) is denoted by e. Perhaps surprisingly, r plays a relatively minor role. The key parameter, as the examples below will indicate, is the sign of $v - e\alpha$.

We call (R, H) *dense* if $v - e\alpha < 0$ and *sparse* if $v - e\alpha > 0$. The irrationality of α comes in at this point, making this a strict dichotomy. We further call (R, H) *rigid* if for all S with $R \subseteq S \subset V(H)$ the rooted graph (S, H) is dense. (As S may be R itself, rigid implies dense.) We call (R, H) *safe* if for all S with $R \subset S \subseteq V(H)$ the rooted graph $(R, H|_S)$ is sparse.

(Here $H|_S$ is the restriction of H to S, simply throw all other vertices away. As S may be $V(H)$ itself safe implies sparse.) Very roughly we think of rigid as meaning dense through and through and safe as meaning sparse through and through. We call $(R, H|_S)$ a subextension of (R, H) and we call (S, H) a nailextension (we are nailing down some more roots) of (R, H).

Lets look at several examples with $\alpha = \pi/7 = 0.448\cdots$. We select this α because it seems to have no special properties whatsoever.

• Every two vertices have a neighbor. H has y_1 adjacent to x_1, x_2. $r = 2, v = 1, e = 2$ so $v - e\alpha > 0$ and (R, H) is sparse and safe.

• Every three vertices have a neighbor. H has y_1 adjacent to x_1, x_2, x_3, $r = 3, v = 1, e = 3$, and $v - e\alpha < 0$ and (R, H) is dense and rigid.

• Every vertex lies in a K_5. H has y_1, y_2, y_3, y_4, x_1 with all ten adjacencies, $r = 1, v = 4, e = 10$ and $v - e\alpha < 0$ and (R, H) is dense and rigid.

• Every vertex lies in a K_4. H has y_1, y_2, y_3, x_1 with all six adjacencies, $r = 1, v = 3, e = 10$ and $v - e\alpha > 0$ and (R, H) is dense and rigid.

• Every two vertices lie in a K_4 except possibly they are nonadjacent. H has y_1, y_2, x_1, x_2 with five adjacencies (not x_1, x_2) , $r = 2, v = 2, e = 5$, $v - e\alpha < 0$, (R, H) is dense and rigid.

• Every three vertices have a common neighbor which itself has a (different) neighbor. H has y_1 adjacent to x_1, x_2, x_3 and y_2 adjacent to y_1. Here $r = 3, v = 2, e = 4$ and $v - e\alpha > 0$ so that (R, H) is sparse. But (R, H) is not safe since the subextension "every three vertices have a common neighbor" $(S = \{x_1, x_2, x_3, y_1\})$ is not sparse.

• Every four vertices have a common neighbor which itself has a (different) neighbor. H has y_1 adjacent to x_1, x_2, x_3, x_4 and y_2 adjacent to y_1. Here $r = 4, v = 2, e = 5$ and $v - e\alpha < 0$ so that (R, H) is dense. But nailing down y_1, setting $S = R \cup \{y_1\}$, gives (S, H) with $r = 5, v = 1, e = 1$ and $v - e\alpha > 0$, so that y_2 is flapping in the wind and (R, H) is not rigid.

It can be shown that $Ext(R, H)$ holds almost surely if and only if (R, H) is safe. Let us see the intuitive justification. Given the $x_1, \ldots, x_r$ we have $\sim cn^v$ choices for $y_1, \ldots, y_v$ and each choice will have the needed e adjacencies with probabiity p^e, hence the expected number of extensions is $\sim cn^v p^e \sim cn^{v-e\alpha}$. When $v - e\alpha < 0$ this expected number goes to zero so almost surely a random $x_1, \ldots, x_r$ will not have an extension. If there is a subextension $(R, H|_S)$ which is not sparse (and hence dense) almost surely a random $x_1, \ldots, x_r$ can not be extended to $H|_S$ and hence not to H. The converse requires more work.

What about rigid? It is not the case that every three vertices have a common neighbor, indeed a random three vertices almost surely will not have a common neighbor. But *some* sets of three vertices do have a common neighbor. (Take a vertex y_1, take three of its neighbors x_1, x_2, x_3 - those three vertices have the common neighbor y_1.) When x_1, x_2, x_3 have a common neighbor that is a special property of the triple. Its not special when x_1, x_2 have a common neighbor since every pair of vertices have a common neighbor. It will turn out that all special properties of bounded sets of vertices are describable in terms of rigid extensions.

4 Closure

Fix $\alpha \in (0,1)$ irrational and $t \geq 1$. Let G be any graph though we'll be interested in $G \sim G(n,p)$ with $p = n^{-\alpha}$. Let X be any set of vertices of G. We define the t-closure of X, denoted by $cl_t(X)$.

Our first definition of $cl_t(X)$ is algorithmic. We say $y_1, \ldots, y_v$ form an (R, H) extension over $x_1, \ldots, x_r$ if they have the required adjacencies of H between the y_i, y_j and the x_i, y_j. We say $y_1, \ldots, y_v$ forms a rigid extension over $x_1, \ldots, x_r$ if they form an (R, H) extension for some rigid (R, H). Now begin with X. If any $y_1, \ldots, y_v$ with (critically) $v \leq t$ form a rigid extension over X then add those vertices to X. Iterate until there are no further rigid extensions. The final set is $cl_t(X)$.

The second definition is that $cl_t(X)$ is the minimal set Z containing X which does not have any rigid extensions of at most t vertices.

Justifying that these two definitions are equivalent and indeed that they are well defined (e. g. that the first doesn't depend on the order in which rigid extensions are added on) requires a series of relatively elementary combinatorial lemmas which we delete. As an example, $cl_4(x_1, x_2)$ might consist of x_1, x_2; y_1, y_2 adjacent to each other and to both x_1, x_2; y_3, y_4, y_5, y_6 forming a K_5 with y_2; and y_7 common neighbor of x_2, y_1, y_5.

Nonexistence Lemma: For every $t \geq 1$ almost surely $cl_t(\emptyset) = \emptyset$ in $G \sim G(n, n^{-\alpha})$.

Proof: When $(\emptyset, H)$ is rigid (or even just dense) it has v vertices and e edges with $v - e\alpha < 0$ so that the expected number of copies of H is $\sim cn^v p^e$ which goes to zero. Hence almost surely there is no copy of H. With t fixed there are only a finite number of such H's to consider so almost surely none of them exist as subgraphs of G.

Let $x_1, \ldots, x_r \in G$, $x'_1, \ldots, x'_r \in G'$. We see that their t-closures are isomorphic, and write $cl_t(x_1, \ldots, x_r) \cong cl_t(x'_1, \ldots, x'_r)$ if there is a graph isomorphism between the t-closures preserving adjacency, nonadjacency and corresponding x_i to x'_i. When H is the restriction of G to $cl_t(x_1, \ldots, x_r)$ we write $cl_t(x_1, \ldots, x_r) \cong H$, but with the additional understanding that the roots $x_1, \ldots, x_r$ are in specified positions in H. For completion we include the case $t = 0$: We define the 0-closure of X to be X and say $cl_0(x_1, \ldots, x_r) \cong cl_0(x'_1, \ldots, x'_r)$ if the map sending x_i to x'_i is a graph isomorphism on these sets of r vertices. Observe that stating $cl_t(x_1, \ldots, x_r) \cong H$ is a first order predicate. In the example of the preceeding paragraph it would consist of stating the existence of the $y_1, \ldots, y_7$ with their appropriate adjacencies and then, for each of the finite list of possible (R, H) rigid extension with $v \leq 4$, the nonexistence of $z_1, \ldots, z_v$ having those adjacencies over $x_1, \ldots, y_7$. A priori the t-closure might be arbitrarily large and the following lemma plays an important role in limiting its possibilities.

Finite Closure Lemma: For all $\alpha \in (0, 1)$, irrational, $r, t \geq 1$ integers there exists K so that in $G \sim G(n, n^{-\alpha})$ almost surely

$$|cl_t(x_1, \ldots, x_r)| < r + K \text{ for all } x_1, \ldots, x_r$$

Proof: We set $\epsilon = \min(e\alpha - v)/v$ over all integers v, e with $v \leq t$ and $v - e\alpha \leq 0$. Note critically the restriction $v \leq t$ allows us to restrict to a finite number of cases and thus the min does exist and (as α is irrational) is positive. We set $K = \lceil r/\epsilon \rceil$.

Suppose the result false and there was $R = \{x_1, \ldots, x_r\}$ with a larger t-closure. Then there would be a sequence $R = R_0 \subset R_1 \subset \ldots \subset R_l$ with each R_{i+1} rigid over R_i with fewer than t nonroots and R_j having size in $[r + K, r + K + t)$. (That is, continue taking rigid extensions and stop when at least $r + K$ vertices are in the set.) Let H_i be the restriction of G to R_i and set H equal the final H_l. Let (R_{i-1}, H_i) have parameters v_i, e_i. Then H has $V = r + \sum_{i=1}^{l} v_i$ vertices and at least $E = \sum_{i=1}^{l} e_i$ edges. Roughly the r roots are our capital and each extension costs us $e\alpha - v$. Formally

$$V - E\alpha \leq r + \sum_{i=1}^{l} (v_i - e_i)\alpha \leq r - \epsilon \sum_{i=1}^{l} v_i \leq r - K\epsilon < 0$$

The existence of such H would then violate the Nonexistence Lemma.

5 The Almost Sure Theory

To describe the almost sure theory $T = T_\alpha$ we require one more somewhat technical point. When (R, H) is safe we want that every $x_1, \ldots, x_r$ should have an (R, H) extension $y_1, \ldots, y_v$. But we further need that these ys have no additional properties relative to the xs. We define this in the first order world via rigid extensions. Roughly we want to say that any rigid extension over the xs and ys is really just over the xs.

Definition: We say $y_1, \ldots, y_v$ is t-generic over $x_1, \ldots, x_r$ if the following holds: Consider any $z_1, \ldots, z_w$ distinct from the xs and ys with (critically) $w \leq t$ which forms a rigid extension over $x_1, \ldots, x_r, y_1, \ldots, y_v$. Then there are no edges between any z_i and any y_j.

The almost sure theory T_α consists of two schema.

• Nonexistence. For H with v vertices, e edges and $v - e\alpha < 0$: There does not exist a copy of H. To express it in slicker form - for all $t \geq 1$: $cl_t(\emptyset) = \emptyset$.

• Generic Extension. For (R, H) safe, $t \geq 0$. For all $x_1, \ldots, x_r$ there exist $y_1, \ldots, y_v$ such that

1. $y_1, \ldots, y_v$ forms an (R, H) extension over $x_1, \ldots, x_v$.
2. There are no additional edges of the form y_i, y_j or y_i, x_j except those mandated by H.
3. $y_1, \ldots, y_v$ is t-generic over $x_1, \ldots x_v$. (For $t = 0$ exclude this condition.)

We've seen by the Nonexistence Lemma that the A in the Nonexistence schema hold almost surely. We indicate the argument for Generic Extension. Let (R, H) be safe. For any $\vec{x} = (x_1, \ldots, x_r)$ let $N(\vec{x})$ denote the number of (R, H) extension $\vec{y} = (y_1, \ldots, y_v)$. Let $x_1, \ldots, x_r$ be selected randomly so that $N = N(\vec{x})$ becomes a random variable. We have seen that the expectation $\mu := E[N] \sim cn^v p^e$ which goes to infinity like a positive power of n. At heart (and the one fairly technical part of the probability analysis) is a Large Deviation result: For any fixed $\epsilon > 0$

$$\Pr[|N(\vec{x}) - \mu| > \epsilon\mu] = o(n^{-r})$$

Actually the probability can be bounded by $\exp[n^{-\lambda}]$ for a positive λ but the above suffices for our purposes. Here N counts extensions and so is the sum of $\sim cn^v$ indicator random variables (one for each distinct extension) each of which are one (i. e. , the extension is there) with probability p^e. If we could think of N as the binomial distribution with parameters cn^v, p^e then the above large deviation result would follow from standard probability results, known as the Chernoff bounds. The difficulty arises in that the indicator random variables are not independent, the potential extensions have a complex overlap pattern. Most of the potential extensions (as v is fixed and $n \to \infty$) do not overlap and so their indicator random variables are independent. Still, it requires some technical skill, which we omit from this presentation, to show the large deviation result.

Given the Large Deviation result we easily deduce a Counting Theorem: Almost surely the number of extensions $N(\vec{x})$ lies between $\mu(1 \pm \epsilon)$ for *all* choices of $\vec{x}$. This follows since there are only $O(n^r)$ choices for the roots and the failure probability is $o(n^{-r})$ for any particular choice. Now, modulo some combinatorial work, we can deduce Generic Extension. For each $\vec{x}$ the number of (R, H) extensions is $\Theta(n^{v-e\alpha})$. How many of these are not t-generic. There are only a finite number of ways $\vec{y}$ can be not t-generic over $\vec{x}$. One shows that for each such possibility the number of such extensions is (using the Counting Theorem upper bound) at most $O(n^{v'-\alpha e'})$ where $v' - \alpha e'$ is smaller than $v - e\alpha$. Roughly, the existence of a rigid extension would add v_1 vertices and e_1 edges with $v_1 - e_1\alpha < 0$ and that would decrease $v - e\alpha$. Then the total number of non t-generic extensions over $\vec{x}$ is bounded by a constant times a smaller power of n. For n sufficiently large this is smaller than the total number of extensions and therefore some (R, H) extension - indeed, almost all such extensions -will be t-generic.

The completeness of T_α is shown via the Ehrenfeucht game but requires a surprisingly subtle strategy for the Duplicator. Let G, G' be models of T_α, fix the number of rounds $u \geq 1$, and consider the Ehrenfeucht game $\mathrm{EHR}(G_1, G_2; u)$.

Define integers $t_0, t_1, \ldots, t_u$ as follows. Set $t_0 = 0$ and (for convenience) $t_1 = 1$. Given t_i select t_{i+1} with

1. $t_{i+1} \geq t_i$
2. Almost surely in $G(n, n^{-\alpha})$ for every X of size $i+1$ the t_i-closure of X has size at most t_{i+1} vertices outside of X.

Of course, the existence of t_{i+1} requires the Finite Closure Lemma. Now we describe Duplicator's strategy. Let x_j, x'_j denote the vertices of G, G' respectively selected in the j-th round. Let $0 \leq i \leq u$ and set $s = u - i$ for convenience. Duplicator plays so that after the s-th round (equivalently, with i rounds remaining) the t_i-closure of $(x_1, \ldots, x_s)$ and the t_i-closure of $(x'_1, \ldots, x'_i)$ are isomorphic, the isomorphism sending x_i to x'_i.

At the start of the game, setting $t = t_u$, the Nonexistence Schema assures that $cl_t(\emptyset)$ is the same in G and G' so Duplicator is fine. At the end of the game the 0-closures are isomorphic which is precisely the condition for

Duplicator to have won. It thus suffices to show (the hard part) that if this condition is satisfied for i then regardless of Spoiler's move Duplicator has a response that preserves the condition for $i-1$.

To avoid subscripts let us fix i and write $BIG := t_i$, $SMALL := t_{i-1}$, $\vec{x} = (x_1, \ldots, x_s)$, $\vec{x'} = (x'_1, \ldots, x'_s)$. By symmetry we can assure Spoiler plays next in G, let y denote his next move. There are two basic cases that we dub Inside and Outside.

We say y is Inside if $y \in cl_{BIG}(\vec{x})$. As $SMALL \leq BIG$ this then determines $cl_{SMALL}(\vec{x}, y)$ which lies entirely inside $cl_{BIG}(\vec{x})$. Duplicator checks the isomorphism between the BIG-closures of $\vec{x}, \vec{x'}$ and selects y' the vertex corresponding to y under the isomorphism.

Otherwise, y is Outside. Let OLD denote the BIG-closure of $\vec{x}$. Duplicator calculates $cl_{SMALL}(\vec{x}, y)$ and sets NEW equal those vertices of it which aren't already in OLD. Our definition of BIG, which in turn depended on the Finite Closure Lemma, assures us that NEW has at most BIG vertices. Say NEW over OLD forms an (R, H) extension. We need now a combinatorial lemma (proof omitted) that any nonsafe extension contains a rigid subextension. From this it follows that (R, H) must be safe since otherwise there would be a nonempty NEW^- rigid over OLD but then it would be in OLD by the closure definition. Duplicator then goes over to G' and by t-generic extension ($t = SMALL$) finds a NEW' over $OLD' = cl_{BIG}(\vec{x'})$ with precisely the same edges and selects y' the vertex of NEW' corresponding to y. This immediately gives that the $SMALL$-closure of $\vec{x'}, y'$ contains a copy of the $SMALL$-closure of $\vec{x}, y$ and some combinatorial lemmas involving t-genericity insure that it contains nothing more and that the two $SMALL$-closures are isomorphic.

This shows that T_α is complete and hence the Zero-One Law.

6 The Case p Constant

One of the original motivations for considering this area was a beautiful result shown independently by Glebskii et. al. and Fagin. Let $0 < p < 1$ be constant. They then showed a Zero-One Law for $G(n, p)$, that every first order A holds either almost surely or almost never.

With our machinery the proof is quite quick. The theory T is given by one schema.

(For all $r, s \geq 0$:) For all distinct $x_1, \ldots, x_r, y_1, \ldots, y_s$ there exists a distinct z adjacent to all of the x_i and to none of the y_j.

Fix r, s, p. Call z a witness (relative to the xs and ys) if it has precisely the desired adjacencies. Each z has probability $\epsilon := p^r(1-p)^s$ of being a witness. The events of being a witness are independent (involving disjoint edgesets) so the probability is $(1-\epsilon)^{n-r-s}$ that there is no witness. There are $\binom{n}{r}\binom{n-r}{s} \leq n^{r+s}$ choices for the xs and ys. Hence the probability that any such choice produces no witness is $\leq n^{r+s}(1-\epsilon)^{n-r-s}$. Fixing r, s, p fixes $\epsilon > 0$ and exponential decay kills off polynomial growth so the failure probability goes to zero.

The graphs G modelling T are said by Peter Winkler to have the Alice's Restaurant property. Members of a certain generation may remember the refrain: You can get anything you want at Alice's Restaurant. All possible witnesses are there.

Let G, G' model T. Duplicator's stategy is simplicity itself. Staying alive. When x_i is played in G Duplicator looks for $x_i' \in G'$ with the appropriate adjacencies to the previously selected vertices. By the Alice's Restaurant property she never gets stuck.

7 Countable Models

Whenever we have a Zero-One Law we have the complete theory T of those sentences holding almost surely. By the Gödel Completeness Theorem such a theory must have a finite or countable model. The models cannot be finite since for every $r \geq 1$ the sentence "There exist distinct $x_1, \ldots, x_r$" is in the almost sure theory since it holds for all $n \geq r$. Thus T must have a countable model - in our case a countable graph G. What does G look like? The first question is whether G is unique - that is, whether T is $\aleph_0$-categorical.

Consider first the Alice's Restaurant theory T for p constant. This is $\aleph_0$-categorical by an elegant argument. Let G, G' be two countable models of T, both labelled by the positive integers. We build up an isomorphism $\Phi : G \to G'$ by alternating Left Stages and Right Stages. After n steps the map Φ will map n elements of G into n elements of G' preserving adjacency and nonadjacency. For a Left Stage let x be the least element of G for which $\Phi(x)$ is not defined. We require of $\Phi(x)$ that for any $a \in G$ for which $\Phi(a)$ has been defined we want $\Phi(x)$ to be either adjacent or nonadjacent to $\Phi(a)$ depending on whether x is adjacent or nonadjacent to a. By Alice's Restaurant we can find such an x'. In the Right Stage we reverse the roles of G, G'. Let x' be the least element of G' for which $\Phi^{-1}(x')$ is not defined and find $x = \Phi^{-1}(x')$ with the appropriate adjacencies. By step $2n$ vertices

$1, \ldots, n$ have been used up in both G and G' so that at the end of this infinite process all vertices have been used up and Φ is a bijection giving the desired isomorphism. The countable graph G satisfying Alice's Restaurant is sometimes called the Rado graph in honor of the late Richard Rado.

What about the theory T_α for $0 < \alpha < 1$ irrational. This is not $\aleph_0$-categorical. We indicate two arguments that create (well, prove the existence) of different countable models.

Consider rigid extensions with $r = 1$, so of the form $(\{x\}, H)$, with parameters v, e where $(\emptyset, H)$ is safe. (With $\alpha = \pi/7$ an example is $H = K_5$.) For such H almost surely there exist copies of H but most vertices do not lie in such copies. Suppose $(\{x\}, H_i)$ is an sequence of such extensions with parameters v_i, e_i. For any s define the graph H^s to be the of $H_1, \ldots, H_s$ considered as disjoint vertex sets except for the common vertex x. Suppose further that there almost surely exists a copy of H^s. Such a sequence can be shown to exist for any α by employing a little number theory. The key is to find v_i, e_i such that $v_i - e_i\alpha$ is only very slightly negative. Now we can create a model in which some element is in a copy of H^s for all s. We add a constant symbol c to our logic and add the infinite schema (for $s \geq 1$) that c is in a copy of H^s. Any finite segment of this system is consistent since in T itself one has that there exists a copy of H^s. By compactness there exists a model and the element corresponding to c has the desired property.

Now we create a special countable graph G_α that models T_α. The vertices will be the positive integers. For every safe rooted graph (R, H) and every $r = |R|$ distinct integers $\vec{x} = (x_1, \ldots, x_r)$ consider the *witness demand* that there must exist $\vec{y} = (y_1, \ldots, y_v)$ forming an (R, H) extension over $\vec{x}$. Witness demands would include, continuing with our standard $\alpha = \pi/7$ example, that there exists y_1 adjacent to $167, 233$ or that there exist y_1, y_2, y_3 forming a K_4 with 26. We include the case $R = \emptyset$ so that one demand is that there exist y_1, y_2 forming an edge. Turn the witness demands into a countable list. Now satisfy them one by one using new points in a minimal way . That is, when we need y_1 adjacent to $167, 233$ pick a vertex, say 23801 that has not been touched before (at any stage only a finite number of points have been touched) and join it to $167, 233$ and nothing else. There are two very nice properties of this construction. First G_α is a model of T_α. (As you might expect these minimal extensions are t-generic for all t.) Second, and quite surpisingly, G_α is unique. That is, it does not depend on the ordering of the witness demands nor on the choice of new points to satisfy them. These graphs G_α seem quite intriguing objects worthy of study simply as countable graphs. For any finite set X of vertices let us define the closure

$cl(X)$ as the union of the t-closures of X over all t, noting this is not a first order concept. In this procedure at some finite time all vertices of X have been touched. Let Y be the value of $cl(X)$ at that moment. After this time all extensions of Y are via safe extensions and one can show that $cl(X)$ remains the same. That is, in G_α all finite sets have finite closure.

The two models created are different since in the first there is an x with $cl(\{x\})$ infinite while in the second there is no such x.

8 A Dynamic View

We have seen that for fixed irrational $\alpha \in (0,1)$ any first order A holds almost surely or almost never in $G(n, n^{-\alpha})$. Now we consider A fixed and vary α - thinking roughly of the evolution of the random graph as we consider $p = n^{-\alpha}$ with α decreasing from one to zero. To study that evolution we define

$$f_A(\alpha) = \lim_{n\to\infty} \Pr[G(n, n^{-\alpha}) \models A]$$

To avoid the problems at rational α we simply define the domain of f_A to be the irrational $\alpha \in (0,1)$. Our goal is to describe the possible functions f_A. Note that $f_A(\alpha) = 1$ when A is in the theory T_α, otherwise $f_A(\alpha) = 0$. We have given an explicit description of the theories T_α. In this sense the function f_A is described independently of probabilistic calculation. We seek to understand the relationships between the continuum of theories T_α.

We begin with a continuity result. Fix A and irrational α. We claim that $f_A(\beta)$ is constant in some interval $(\alpha-\epsilon, \alpha+\epsilon)$ around α. Suppose A is in T_α (otherwise take $\neg A$). Then A follows from a finite number of axioms of T_α. These in turn depend on notions of dense and sparse rooted graphs which depend on whether $v - e\alpha$ is positive or negative. For any particular v, e whatever the sign of $v - e\alpha$ that sign remains constant in some interval around α. The finite number of axioms leads to a finite number of pairs v, e and so all signs remain constant in some interval. For β in that interval T_β has these same axioms and so A is in T_β. (It is known, however, that the theories T_α are all different. Between any two α, α' lies a rational a/b and it is known that there is a graph H such that the existence of a copy of H has threshold function $n^{-a/b}$.)

The discontinuities of f_A must therefore come at the rational $a/b \in (0,1)$. We define the spectrum $Sp(A)$ to be those rational points of discontinuity. The classical theory of Random Graphs gives natural examples. Existence of a K_4 has spectrum $\{2/3\}$. Existence of a K_5 has spectrum $\{1/2\}$. We can put these together: "There exists a K_4 and there does not exist a K_5" to give spectrum $\{2/3, 1/2\}$ - here as G evolves $\Pr[A]$ starts near zero, jumps to one at $n^{-2/3}$ when K_4 appear and back down to zero at $n^{-1/2}$ when K_5 appear. With some technical work it is not difficult to get any finite set of rationals in $(0, 1)$ as a spectrum in this way. This author once conjectured that all spectra were such finite sets. That proved not to be the case.

9 Infinite Spectra via Almost Sure Encoding

Here we will describe a first order A with an infinite spectrum. The central idea will be to take a second order sentence and give it an almost sure encoding in the first order language.

For definiteness we will work near $\alpha = \frac{1}{3}$. By a $K_{3,k}$ is meant a set $x_1, x_2, x_3; y_1, \dots, y_k$ with all y_j adjacent to all three xs. Basic random graph theory gives that the sentence "There exists a $K_{3,k}$" has threshold function $n^{-1/3-1/k}$. (There are $e = 3k$ edges and $v = 3 + k$ vertices and $(\emptyset, K_{3,k})$ is sparse and safe if and only if $v - e\alpha > 0$.) Let $N(x_1, x_2, x_3)$ denote the set of common neighbors of x_1, x_2, x_3. Then for $\frac{1}{3} + \frac{1}{k} > \alpha > \frac{1}{3} + \frac{1}{k+1}$ the maximal size $|N(x_1, x_2, x_3)|$ is k. Consider then the property, call it A^*, that the maximal size $|N(x_1, x_2, x_3)|$ is even. This would have all values $\frac{1}{3} + \frac{1}{k}$ as spectral points. It is not possible to write this property in the first order language. We shall, however, give an almost sure encoding, a first order sentence that almost surely has the same truth value as A^*.

Lets look in the second order world. How can we say that a set S (which will be $N(x_1, x_2, x_3)$ in our application) has even size. We write:

$$EVEN(S) : \exists_R \forall_x \neg R(x, x) \wedge \forall_{x,y} R(x, y) \leftrightarrow R(y, x) \wedge \forall_{x \in S} \exists!_{y \in S} R(x, y)$$

That is, there exists an areflexive symmetric binary relation on S (i. e. a graph) which is a matching - each vertex has precisely one neighbor. How can we say that S is bigger (or equal) in size to T. Similarly we write $BIGGER(S, T)$ that there exists an areflexive symmetric binary relation R that yields an injection from $T - S$ to $S - T$. For every $y \in T - S$ there is a $x \in S - T$ with $R(y, x)$ and we do not have $R(y_1, x)$ and $R(y_2, x)$ for distinct $y_1, y_2 \in T - S$ and $x \in S - T$. Now we can write A^* in second order:

$$A^* : \exists_{x_1,x_2,x_3} EVEN[N(x_1, x_2, x_3)] \wedge$$

$$\wedge \forall_{z_1,z_2,z_3} BIGGER[N(x_1, x_2, x_3), N(z_1, z_2, z_3)]$$

Now for the almost sure encoding. Define the first order ternary predicate (considering u as a variable symbol)

$$R_u(x,y) := \exists_v [v \sim x \wedge v \sim y \wedge v \sim u],$$

that u, x, y have a common neighbor. Our basic (though it will need modification) idea is to replace the second order $\exists_R$ with the first order $\exists_u$ and then to replace all instances of the binary R with the now binary R_u.

Representation Lemma: For any s and any symmetric areflexive R on $1, \ldots, s$ that holds for l pairs with $l < \frac{k}{3}$

$$\forall_{x_1,\ldots,x_s} \exists_u \bigwedge_{1 \le i < j \le s} (R_u(x_i, x_j) \leftrightarrow R(i,j))$$

is a theorem of $T = T_\alpha$ for all $\frac{1}{3} + \frac{1}{k} > \alpha > \frac{1}{3} + \frac{1}{k+1}$.

Consider the rooted graph, call it (S, H) with roots $1, \ldots, s$, nonroot u, and then for each $1 \le i < j \le s$ nonroot v_{ij} with edges from v_{ij} to i, j, u. (S, H) has $v = 1 + l$ nonroots and $e = 3l$ edges. Our bound on l assures that $v - e\alpha > 0$ so that (S, H) is sparse, and some easy combinatorial work shows that it is safe as well. In T_α we have the 1-Generic Extension axiom for (S, H). For all $x_1, \ldots, x_s$ there exists u and the v_{ij} having the above edges and no more so that when $R(i,j)$ we do have $R_u(x_i, x_j)$. Suppose now $\neg R(i,j)$, can u, x_i, x_j have a common neighbor? A common neighbor to three vertices is a rigid extension in our range $\alpha > \frac{1}{3}$ so this would violate 1-genericity.

We outline a second argument more for those in random graphs. Set $p = n^{-1/3-\epsilon}$ so that $\frac{1}{k} > \epsilon > \frac{1}{k+1}$. Any particular $R_u(x,y)$ holds with probability roughly $np^3 \sim n^{-3\epsilon}$, that being the expected number of common neighbors. Say u is a witness if $R_u(x,y)$ holds for the l needed pairs. Then u would be a witness with probability roughly $n^{-3l\epsilon}$. There are n potential witnesses so the expected number of witnesses would be roughly $n^{1-3l\epsilon}$. As $3l\epsilon < 1$ this expected number goes to infinity and almost surely for every choice of the xs there is one. There are a number of questions here (for one thing, u, u' being witnesses are no longer fully independent events) that need to be fleshed out but this can be turned into a full proof.

We have a small technical problem. We want to say $EVEN(S)$ where $S = N(x_1, x_2, x_3)$ has at most k elements by saying there is a matching R. Such an R would have perhaps $k/2$ edges while our representation lemma only gives us R_u with at most $k/3$ edges. We puff up the representation lemma by replacing $\exists_R$ with $\exists_{u_1,u_2}$ and replacing R with $R_{u_1} \vee R_{u_2}$. Now we represent all R with up to just less than $2k/3$ edges. To write it out in full, "$N(x_1, x_2, x_3)$ is even" is replaced by

There exist u_1, u_2 such that for all y adjacent to x_1, x_2, x_3 there exists a unique $y' \neq y$ adjacent to x_1, x_2, x_3 with either y, y', u_1 or y, y', u_2 having a common neighbor.

Similarly, $BIGGER(S,T)$ may require an injection R of k edges. We therefore replace $\exists_R$ with $\exists_{u_1,u_2,u_3,u_4}$ and R with $R_{u_1} \vee R_{u_2} \vee R_{u_3} \vee R_{u_4}$. With this $BIGGER(N(x_1, x_2, x_3), N(x_1' x_2' x_3'))$ becomes a first order predicate. We have given an almost sure encoding that transforms second order A^* into a totally first order [though hardly natural to those in graph theory!] sentence A which has the desired infinite spectrum.

The notion of almost sure encoding is an intriguing one and will appear several more times. One is given a property P in some large language L^+ and one wishes to find (or, in one example later, to disprove the existence of) a sentence A in a given smaller language L which is an almost sure encoding of it. By this we mean that the probability of P, A differing in truth value goes to zero as the model size goes to infinity. Of course, one also has to fix the probability measure, in our case $G(n, p(n))$ with some particular $p(n)$. Hella, Kolaitis and Luosto have called two languages L, L' almost everywhere equivalent if for every P in one language there is an A in the other where, as above, the probability of P, A differing in truth value goes to zero as the model size goes to infinity. One particularly intriguing problem they give involves $G(n,p)$ with $p = \frac{1}{2}$: Is monadic existential second order logic almost everywhere equivalent to monadic universal second order logic? They conjecture that the answer is no but it does seem difficult to show negative results about the existence of an almost sure encoding.

10 The Jump Condition

We have already mentioned that the theories T_α are all distinct. However, if we fix the quantifier depth u of the sentences we are examining then the fall into definite intervals. Lets recall the sequence $t_0, \ldots, t_u$ from § 5. We had $t_0 = 0, t_1 = 1$ and $t_{i+1} = \max[t_i, \lceil (u-i)\epsilon^{-1} \rceil]$ where ϵ was the minimum value of $v^{-1}(e\alpha - v)$ over all integers v, e with $v \leq t_i$ and $v - e\alpha \leq 0$. We try to define this sequence for rational α as well. It doesn't always work. Take, for example, $u = 5$ and $\alpha = \frac{1}{3} + 10^{-6}$. With $t_1 = 1$ we take $v = 1, e = 3$ to give $\epsilon = 3 \cdot 10^{-6}$. This yields a t_2 roughly $\frac{4}{3}10^6$ which is bigger than the numerator of $10^6 + 1$ of α. Now in trying to define t_3 we have v, e with $v \leq t_2$ and $v - e\alpha = 0$ so that $\epsilon = 0$ and the process explodes.

This isn't a surprise, the Zero-One Law isn't supposed to hold for rational α. But it will hold on sentences of quantifier depth u if the rational α is not too rational. To be precise, let XPL_u denote the set of rational α for which the sequence $t_0, \ldots, t_u$ is not well defined together (a technical point) with those α for which the sequence is well defined and α has numerator at most t_u. For $\alpha \notin XPL_u$ we do get a Zero-One Law. It turns out that XPL_u is a well ordered set under the ordering $>$. (There is a lot of pretty number theory involved in studying XPL_u which is quite remniscent of continued fractions. The example above actually shows $\frac{1}{3} + \frac{1}{m} \in EXP_5$ for all large integers m so that EXP_5 is infinite. Here $\frac{1}{3}$ is an accumulation point of EXP_5 but only from larger values.) That is, for every $a/b \in XPL_u$ (except the smallest) there is an $(a/b)^- \in XPL_u$ which is the biggest element of XPL_u smaller than a/b. Then XPL_u splits the unit interval into intervals I from (going down) a/b to $(a/b)^-$. (We include the I from the smallest value of XPL_u to zero.) Inside each interval the sequences $t_0, \ldots, t_u$ are the same. Further, the truth value of any A of quantifier depth u remains the same as α ranges over such an I. (Basically, one only needs notions of safe and dense rooted graphs up to $v = t_u$ and these notions are the same for all α in the interval.) To rewrite as a condition on possible f_A:
Jump Condition: If $f = f_A$ for some first order A then there is a u such that f is constant on each interval I defined by the splitting set XPL_u.

11 The Complexity Condition

For $\alpha \in (0, 1]$ rational let us define $g_A(\alpha)$ to be the limiting value of $f_A(\alpha - \epsilon)$ as ϵ approaches zero from above. Since EXP_u is well ordered under $>$ this is well defined. Indeed, for $\alpha \in EXP_u$ this gives the value of f_A on the interval from α to the next α^-. Since the intervals I defined above partition the unit interval g_A will determine f_A.

For $\alpha \in (0, 1]$ we define a theory T_α^-. This will be the limiting theory of the $T_{\alpha+\epsilon}$ as ϵ approaches zero from above. Recall that the splitting into dense and sparse rooted graphs was not a strict dichotomy for α rational because of the possibility that $v - e\alpha = 0$. In T_α^- we simply consider such rooted graphs as sparse, as that is their status in $T_{\alpha+\epsilon}$ with ϵ positive. This can be shown to give a complete theory and $g_A(\alpha) = 1$ precisely when A lies in this theory.

We have a most surprising complexity condition on the functions g_A.
Complexity Condition:

$$\{0^a 1^b : A \in T^-_{a/b}\} \in PH$$

To see this, let us fix the quantifier depth u and consider how difficult it is to find if $A \in T^-_{a/b}$ as a function of the denominator b. We can as before define the sequence $t_0, \ldots, t_u$. Here having defined t_i we define ϵ by only looking at those v, e with $v \leq t_i$ and $v - e(a/b)$ strictly negative. But then $v - e(a/b)$ has denominator at most $t_i b$ and so $\epsilon \geq (t_i b)^{-1}$. Other terms (considering u fixed) supply bounded factors, basically t_i goes up by at most a factor of b as i increases. That is, $t_i = O(b^i)$.

We can write any A of quantifier depth u in the form

$$A : Q_{x_1} Q_{x_2} \cdots Q_{xu} P(x_1, \ldots, x_u)$$

where Q is either $\exists$ or $\forall$, very possibly taking different values at different times, and P is a Boolean expression of the atoms $x_i = x_j$ and $x_i \sim x_j$. The truth value of A in $T^-_{a/b}$ can now be turned into a game between two players. We'll call them Spoiler and Duplicator as before, though this game is not the Ehrenfeucht Game. Duplicator's object is to show A is a consequence of $T^-_{a/b}$, Spoiler tries to show it is not.

`The Game Board.` The game board has levels $0, 1, \ldots, u$. Each level has a finite set of positions. At level 0 are the possible values of $cl_0(x_1, \ldots, x_u)$. [Recall that these are determined by the graph on $\{x_1, \ldots, x_u\}$ and, to be formally correct, the equalities amongst the x_i.] At level i are the possible values of the t_i-closure of $x_1, \ldots, x_{u-i}$. When $i = u$, the top level, there is only one possible t_u-closure of $\emptyset$, namely $\emptyset$ so there is only a single position.

`The Initial Position.` The top level position $\emptyset$.

`The Winning Final Positions.` The 0-position determines the truth value of $P(x_1, \ldots, x_u)$ - call a 0-position winning if P is true, otherwise losing.

`The Permitted Moves.` All moves go down one level. Let H, H' be positions on the i and $i-1$ level respectively. Moving from H to H' is permitted if and only if in $T^-_{a/b}$ the following is a theorem: Given any $x_1, \ldots, x_{u-i}$ with t_i-closure H there exists x_{u-i+1} such that the t_{i-1}-closure of $x_1, \ldots, x_{u-i}, x_{u-i+1}$ is H'. We had argued that the T_α are complete via the Ehrenfeucht game but it could have been done syntactically. The key result is that in T_α for any positions H, H' on the $i, i-1$ level either the above is a theorem or it is a theorem that: Given any $x_1, \ldots, x_{u-i}$ with t_i-closure H

there *does not* exists x_{u-i+1} such that the t_{i-1}-closure of $x_1, \ldots, x_{u-i}, x_{u-i+1}$ is H'.

`The Rules of the Game.` There are u rounds. On the i-th round when x_i is quantified existentially (i. e. $Q = \exists$) it is Duplicator's move, when it is quantified universally it is Spoiler's move. In either case the permitted moves are given above so that the position moves through the levels and at the end of the u rounds is on the bottom level. Those positions have been designated winning and losing, and Duplicator wins or loses accordingly.

This game description works for any T_α or $T^-_{a/b}$. But with $T^-_{a/b}$ we can bound the game complexity by noting that each position is given by a graph (together with designated vertices) of size polynomial in b, certainly $O(b^u)$, and hence can be described by a sequence of bits of length $O(b^{2u})$. Therefore winning the game has complexity in the Polynomial Heirarchy at level u.

Well, not quite. We also have to examine whether a move H to H' is permissible. To "prove" that the move is permissible Duplicator draws the picture of H and H'. When the move is Inside she simply designated the new move x_{u-i+1} and the set H' which is the new closure. When the move is Outside she gives which vertices of H are still in H' plus adds the new vertices (called NEW in the completeness proof) with all edges and designated vertex x_{u-i+1}. She further lists the sequence of rigid extensions that give the t_{i+1}-closure. All this can be done with a polynomial length string. Now Spoiler is allowed a polynomial length string to show that Duplicator has been duplicitous. He can show that one of the rigid extensions is not really rigid by nailing down some vertices so that the extension becomes sparse. He can show (in the Outside case) that NEW is not really safe over H by demonstrating a dense subextension. Finally, he can show that the t_{i+1} closure is more that H' by exhibiting, inside Duplicator's picture of $H \cup H'$, a dense extension. (There is a theorem that dense extensions must contain rigid subextensions so he need not show that his extension is rigid.) This shows that the permissibility of a move is in the second level of the Polynomial Heirarchy.

Remarkably, the Jump Condition and the Complexity Condition characterize the possible functions f_A. We have seen, albeit in outline form, that these conditions are necessary. That there are sufficient is technically quite challenging. This result is due to Gabor Tardos.

12 Nonconvergence via Almost Sure Encoding

Let us turn to the random *ordered* graph $G_<(n,p)$. The underlying model is still a vertex set Ω of size n and a probability space of graphs on Ω where each pair of vertices is adjacent with independent probability p. In addition, the set Ω is totally ordered by a built-in relation $<$. This relation is part of the language. For convenience we can assume $\Omega = \{1,\ldots,n\}$. Now 1 is uniquely defined as that element with nothing less than it and 2 is uniquely defined as that element with only 1 less than. We can that express $1 \sim 2$ by the first order sentence:

$$\exists_x \exists_y (x \neq y) \wedge (x < y) \wedge [\forall_z z < y \rightarrow z = x] \wedge x \sim y$$

This event (for $n \geq 2$) has probability p. We shall write $y = x+1$ if $x < y$ and there is no z in between them. When $y \neq 1$ we write $x = y-1$ when $y = x+1$. Note, however, that addition and subtraction are in general not defined in this language.

We shall restrict our attention to $p = \frac{1}{2}$. The example above shows that there is no Zero-One Law, that $\Pr[A]$ need not converge to zero nor one. We aim for the following stronger negative result of Compton, Hansen and Shelah.

Theorem: There is an A for which $\lim_{n\to\infty} \Pr[G_<(n,\frac{1}{2}) \models A]$ does not exist.

The central idea is to encode arithmetic on an ordered set S, first using second order language and then in first order with an almost sure encoding. The second order encoding is standard. We say that on S there exist ternary relations $+(x,y,z), *(x,y,z)$ (with the interpretations $x+y=z$ and $x \cdot y = z$ respectively) such that

1. $+(x,1,z)$ if and only if $z = x+1$ as described above.
2. When $y \neq 1$, $+(x,y,z)$ if and only if $+(x,y-1,z-1)$
3. $*(x,1,z)$ if and only if $z = x$
4. When $y \neq 1$, $*(x,y,z)$ if and only if there exists u with $*(x,y-1,u)$ and $+(x,u,z)$

When this occurs we say S is arithmetizable. Now for the almost sure encoding. For $c \leq d$ we write $R_{c,d}(x,y,z)$ if $x,y \leq z$ and (critically) there exists e with $c \leq e < d$ such that e is adjacent to x,y,z and no other elements of S. We say S is first order arithmetizable if there exist c, d and c', d' such that $R_{c,d}, R_{c',d'}$ have the properties of plus and times enumerated above. For our specific purposes we shall consider only S of the form $\{1, \ldots, u\}$ though one could give similar results for more general S with a bit more technical work. We make all logarithms to base 2 in what follows for definiteness.

Representation Lemma: Let $u \leq 0.9 \log^{1/3} n$ Then almost surely there exist $c \leq d$ such that $R_{c,d}$ is the ternary relation $+$ on $\{1, \ldots, u\}$ and also $c \leq d$ such that $R_{c,d}$ is $*$.

Let $+$ have s instances so that $s < u^2$. Consider a pair c, d with $u < c$ and $d = c + s$. Call c a witness if $R_{c,d}$ is indeed $+$ on $\{1, \ldots, u\}$ There is an arrangement (indeed, many such) of the edges between $\{1, \ldots, u\}$ and $\{c, \ldots, d-1\}$ such that c is a witness. This occurs if us pairs have a particular set of adjacencies (and no more) and so has probability 2^{-us} of occurring. There are $\sim n$ potential witnesses c so that the expected number of witnesses is bigger than roughly $n2^{-us}$. We've bounded u so that $us < u^3 < (0.9)^3 \log n$ and so this expected number goes to infinity. Some technical work shows that almost surely there is a witness. [Actually, the technical work isn't so difficult here. We can pick $\sim c'n \log^{-1/3} n$ values c so that the intervals $[c, d)$ are disjoint and so the events that c is a witness are mutually independent over those different c.] Representing $*$ is the same. Indeed, with further technical work (perhaps modifying the bound on u) one could almost surely represent every ternary, even k-ary, relation R.

Similar arguments, which we exclude, show that when $u > C \log^{1/3} n$ (C a computable absolute constant) than the representation lemma almost surely fails and $\{1, \ldots, u\}$ is not first order arithmetizable. For definiteness let us take $C = 900$. Now the maximal u such that $\{1, \ldots, u\}$ is determined up to a factor of 1000.

Once we have arithmetized $\{1, \ldots, u\}$ we are off to the races. We can say that u is prime, that u is a Fermat prime, there is a large spectra here. Certainly we can talk about $\log u$.

Now we can give our first order sentence A: There exists u such that

1. $\{1, \ldots, u\}$ is first order arithmetizable
2. $\{1, \ldots, u+1\}$ is not first order arithmetizable
3. $\log u$ modulo 40 is one of $1, 2, \ldots, 20$.

Why does this work? The size n of the model almost surely determines u up to a factor of 1000 and so $\log u$ is almost surely determined up to an additive term of 10. For some n this range of $\log u$ will all be in $1, \dots, 20$ modulo 40 while for other n this range will all be in $21, \dots, 39, 0$ modulo 40. This gives infinite subsequences of n on which our sentence has limiting probility one and zero respectively, the worst kind of nonconvergence.

The almost sure encoding can be used to show nonconvergence by encoding arithmetic in other contexts. We examine, in outline form, $G(n, n^{-1/4})$. Note that we do not include $<$ as a built in predicate here. We arithmetize a set S in the second order language by saying that there exists a binary $<$ and ternary $+, *$ with the desired first order properties. For $u \notin S$ we define a ternary R_u on S letting $R_u(x, y, z)$ be the first order property that u, x, y, z have a common neighbor. Also for $u, x \notin S$ we have the binary relation $R_{u,x}(y, z) = R_u(x, y, z)$. [We actually need further technical work here in that such relations are symmetric while $<$ is not.] We say S is first order arithmetizable if there exist u_1, u_2, u_3, u_4 such that $R_{u_1,u_2}, R_{u_3}, R_{u_4}$ play the role of $<, +, *$. At $p = n^{-1/4}$ any four vertices have probability $(1 - p^4)^{n-4} \sim e^{-1}$ of having no common neighbor. Basically, each R_u acts like an independent (this part takes some technical work) random ternary predicate with probability of occurance $1-e^{-1}$. Key here is that both $1-e^{-1}$ and e^{-1} are bounded away from zero. Letting S have size s, a given u witnesses a particular ternary R with probability at least e^{-t} where $t = \binom{s}{3}$ is the number of triples. The expected number of witnesses is at least ne^{-t}. For $s \leq \ln^{1/3} n$ this goes to infinity and one can show that almost surely $+, *, <$ are represented. We cannot quantify over all subsets S in the first order language but instead look at sets $S = N(x_1, x_2, x_3, x_4)$, the set of common neighbors of x_1, x_2, x_3, x_4. One can show that there are such S of all sizes up to roughly $\ln n / \ln \ln n$. On sets S, T of size $O(\ln^{1/3} n)$ we can say $BIGGER(S, T)$ in the first order language (as done in § 9) by saying there exist u_1, u_2 so that R_{u_1,u_2} gives an injection from T to S. It is then a first order property of x_1, x_2, x_3, x_4 that $S = N(x_1, x_2, x_3, x_4)$ is arithmetizable but there is no "bigger" arithmetizable $S' = N(x'_1, x'_2, x'_3, x'_4)$. Such S would almost surely have size $\Theta(\ln^{1/3} n)$. But when S is arithmetizable we can say a wide variety of things about its size u. In particular, we get a nonconvergent sentence by saying that there exist x_1, x_2, x_3, x_4 such that the size $u = |N(x_1, x_2, x_3, x_4)|$ has $\log u$ between 1 and 20 modulo 40.

13 No Almost Sure Representation of Evenness

In this section we restrict ourselves to the random ordered graph $G_<(n,p)$ with $p = \frac{1}{2}$. Set, for any property A,

$$f_A(n) = \Pr[G_<(n,p) \models A]$$

We shall outline the proof of the following result of Saharon Shelah:

Theorem: For any first order A

$$\lim_{n\to\infty} f_A(n+1) - f_A(n) = 0$$

This provides an interesting counterpoint to the Compton, Hansen, Shelah result discussed earlier. There are A for which $f_A(n)$ does not converge but it cannot oscillate back and forth too fast. There is a very nice corollary: There is no first order sentence that provides an almost sure representation for the property that the number n of vertices is even. For such an A would have $f_A(2n) \to 1$ and $f_A(2n+1) \to 0$ which would contradict the slow oscillation of Shelah's Theorem. We find in general that it is quite difficult to prove negative results about almost sure representation and in this context Shelah's result is particularly striking.

We link $G_<(n,p)$ and $G_<(n+1,p)$ be the following procedure. Take a random graph on $2n+1$ ordered vertices, call it $G \sim G_<(2n+1,p)$. Restricting to a random subset S of size precisely n gives $G^{(n)}$, with distribution that of $G_<(n,p)$. Restricting to a random set S of size precisely $n+1$ similarly gives $G^{(n+1)} \sim G_<(n+1,p)$. We thus have

$$f_A(n+1) - f_A(n) = \sum_G \mu(G) \left[\Pr[G^{(n+1)} \models A] - \Pr[G^{(n)} \models A]\right]$$

where $\mu(G)$ is the probability $G_<(2n+1,p)$ is G. Shelah actually shows that for *every* G on $2n+1$ ordered vertices

$$\left|Pr[G^{(n+1)} \models A] - \Pr[G^{(n)} \models A]\right| \to 0$$

Fix G and a property A. Consider the property that G restricted to S satisfies A as a function of S . For example, a sentence such as

$$\exists_x \forall_y \exists_z z \sim y \wedge y \sim x$$

would turn into

$$\exists_x (x \in S) \wedge [\forall_y (y \in S) \to \exists_z (z \in S) \wedge (z \sim x) \wedge (z \sim y)]$$

Such an property A^* is a Boolean function of the variables $x \in S$ for $x = 1, \ldots, 2n+1$. Here we turn to circuit complexity - the function may be represented by a circuit with primitives $x \in S$. Each $\exists_x$ is an OR-gate with fan-in $2n+1$ (that is, all x) and each $\forall_x$ is an AND-gate with fan-in also $2n+1$. The statements $x \sim y$ and $x < y$ then have definite truth values and so do not appear in the circuit. A^* is then represented by a bounded depth polynomial size circuit. It is a deep theorem of circuit complexity (due originally to Razborov) that such a circuit cannot determine majority - that is, cannot be true if and only if at least half of the $2n+1$ inputs are true. Some further technical work shows that no such circuit can distinguish between a random n and $n+1$ inputs being true - that the difference of the probability the circuit yields true in the two experiments must tend to zero. This gives Shelah's result.

14 The Ehrenfeucht Game

The Ehrenfeucht Game is a powerful and very general method for showing that two models have (or do not have) the same first order properties. We consider first the specific example of graphs. Let G, H be two graphs and let t be a positive integer. We describe the Ehrenfeucht Game $\text{EHR}(G, H; t)$.

`The Board:` A copy of G and a copy of H on disjoint vertex sets.

`The Players:` Spoiler and Duplicator.

`The Play:` There are t rounds. On the i-th round Spoiler goes first. He selects either a vertex from G or a vertex from H. Then Duplicator goes. She selects a vertex from the graph that Spoiler did not select from. We let x_i denote the vertex selected from G in the i-th round and y_i the vertex selected from H in the i-th round, regardless of who selected them. We note that Spoiler's choice of which graph to choose from can change from round to round.

`The Winner:` Duplicator wins if and only if the map from x_i to y_i preserves adjacency and equality. That is: x_i, x_j are adjacent in G precisely when y_i, y_j are adjacent in H. Further $x_i = x_j$ precisely when $y_i = y_j$.

We note that when the graphs both have at least t vertices there is no point in Spoiler selecting an x_j equal to a previous x_i as then Duplicator would simply select $y_j = y_i$. Hence we could add the requirement that Spoiler always picks a new vertex. Then Duplicator would also always pick a new vertex.

Theorem: Duplicator wins $\text{EHR}[G, H; t]$ if and only if G, H have the same truth values on all first order sentences of quantifier depth t.

We illustrate this fundamental result with an example. Suppose G has an isolated vertex and H does not. The property $\forall_x \exists_y x \sim y$ has quantifier depth $t = 2$. Spoiler selects the isolated vertex $x_1 \in G$ and Duplicator must select some $y_1 \in H$. As y_1 is not isolated Spoiler moves over to H and selects a $y_2 \in H$ adjacent to y_1. Now Duplicator is stuck, there is no $x_2 \in G$ adjacent to x_1 for her to select.

As an immediate corollary: G, H are elementarily equivalent if and only if Duplicator wins $\mathrm{EHR}[G, H; t]$ for every positive integer t. Note, however, that this is not the same as Duplicator winning a game with an infinite number of moves.

Corollary: Let T be a consistent theory with no finite models. Then T is complete if and only if for every two countable models G, H of T and every positive integer t Duplicator wins $\mathrm{EHR}[G, H; t]$.

If T is complete the models G, H are necessarily elementarily equivalent so that Duplicator wins. If T is not complete there is a sentence A so that $T + A$ and $T + \neg A$ are both consistent and so they have countable models G, H. Letting t be the quantifier depth of A, Spoiler would win $\mathrm{EHR}[G, H; t]$.

Let us generalize to first order languages (we could go even further) with a finite numer of relation symbols R of varying arity. This would include the ordered graph (with $<$ as well as adjacency) or the simple unary language (with only one unary U and equality) of § 1. Let G, H be two models of the language. Then $\mathrm{EHR}[G, H; t]$ is played as described above, with Spoiler and Duplicator selecting $x_1, \ldots, x_t \in G$ and $y_1, \ldots, y_t \in H$. For Duplicator to win she now has to preserve all the relations. That is, let R be any relation symbol of arity, say, l. Then $R(x_{i_1}, \ldots, x_{i_l})$ must have the same truth value as $R(y_{i_1}, \ldots, y_{i_l})$ for every choice of $i_1, \ldots, i_l$ from $1, \ldots, t$.

About the references

Among the other surveys of this area we recommend those of Compton [3], Winkler [26], Lynch [15], and this author [23]. The Ehrenfeucht game was first given in [5]. (It was essentially found in earlier work by Fraisse and is sometimes referred to as the Ehrenfeucht-Fraisse game.) The classic Zero-One law for random graphs with $p = \frac{1}{2}$ (often called the uniform distribution) are due to Glebskii et. al. [8] and Fagin [7]. The classic paper that began the theory of random graphs is by Paul Erdős and Alfred Rényi [6]. The basic text on random graphs is Bollobás [2].

The Zero-One Law for $p = n^{-\alpha}$ appeared first in Shelah, Spencer [17]. An approach using the Ehrenfeucht game is given in Spencer [21]. A syntactic proof of the completeness of the T_α is given in Spencer [22]. An examination of the countable models of T_α is given in Spencer [20]. The Alon, Spencer text [1] also includes some of this material.

In this brief paper we have only examined a few examples of random structures. Among the many others we mention Lynch [14] on unary functions; Shelah and Spencer [18] and StJohn and Spencer [24] on random unary predicates with order (*considerably* different from §1!); Łuczak [11] on random partially ordered sets. Łuczak and Shelah [12] consider an interesting random graph model on vertex set $1, \ldots, n$ where the adjacency probability between i and j depends on $|i-j|$.

While we have here restricted ourselves to first order logic there are a number of papers considering stronger logics. Generally, these give negative results that a Zero-One Law or convergence does not always hold. A nice example is given by Kaufman and Shelah [10], giving a nonconvergent second order sentence on $G(n,p)$ with $p = \frac{1}{2}$. Many such results, including those on the random ordered graph given in the text, can be found in Compton, Henson, Shelah [4]. Shelah [16] shows that on the random ordered graph no first order sentence can almost surely encode the evenness of the model. Hella, Kolaitis and Luosto [9] consider the general problem of almost sure equivalence.

Spencer [19] examines the random graph theory of extension statements in some detail. Łuczak and Spencer [13] use some detailed random graph theory to give a near characterization of those $p = p(n)$ (*not* just those of form $n^{-\alpha}$) for which the Zero-One Law holds. Spencer and Tardos [25] give the necessary conditions on the function $f_A(\alpha)$ defined in the text. The proof of sufficiency by Tardos is in preparation.

References

[1] N. Alon, J. Spencer, The Probabilistic Method, John Wiley, 1991

[2] B. Bollobás, Random Graphs, Academic Press, 1985.

[3] K.J. Compton, $0-1$ Laws in Logic and Combinatorics, in *Algorithms and Order* , I. Rival, ed., NATO ASI series, Kluwer Academic Publishers, Dordrecht, 1988, 353-383

[4] K.J. Compton, C.W. Henson and S. Shelah, Nonconvergence, undecidability and intractability in asymptotics problems, Annals Pure Appl. Logic *36* (1987), 207-224

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

[6] P. Erdős and A. Rényi, On the Evolution of Random Graphs, Mat. Kutató Int. Közl *5* (1960), 17-60

[7] R. Fagin, Probabilities in Finite Models, J. Symbolic Logic *41* (1976), 50-58

[8] Y.V. Glebskii, D.I. Kogan, M.I. Liagonkii and V.A. Talanov, Range and degree of realizability of formulas in the restricted predicate calculua, Cybernetics *5*, 142-154 (Russian original: Kibernetica *5* (1969), 17-27)

[9] L. Hella, , P.G. Kolaitis, K. Luosto, Almost everywhere equivalence of logics in finite model theory. Bulletin of Symbolic Logic *2* (1996), 422-443.

[10] M. Kaufmann and S. Shelah, On random models of finite power and monadic logic, Discrete Math. *54* (1983), 285–293

[11] T. Łuczak, First Order Properties of Random Posets, Order *8* (1991), 291-297

[12] T. Łuczak and S. Shelah, Convergence in Homogeneous Random Graphs, Random Structures & Algorithms *6* (1995), 371-392

[13] T. Łuczak and J. Spencer, When Does the Zero-One Law Hold?, J. Amer. Math. Soc. *4* (1991), 451-468

[14] J. Lynch, Probabilities of first-order sentences about unary functions, Trans. Amer. Math. Soc. *287* (1985), 543-568

[15] J. Lynch, Special Year on Logic and Algorithms Tutorial Notes: Random Finite Models, DIMACS Technical Report 97-56 (1997).

[16] S. Shelah, Very weak zero one law for random graphs with order and random binary functions, Random Structures & Algorithms *9* (1995), 351-358.

[17] S. Shelah and J. Spencer, Zero-One Laws for Sparse Random Graphs, J. Amer. Math. Soc. *1* (1988), 97-115

[18] S. Shelah and J. Spencer, Random Sparse Unary Predicates, Random Structures & Algorithms *5* (1994), 375-394

[19] J. Spencer, Threshold Functions for Extension Statements, J. Comb. Th - Ser A *53* (1990), 286-305

[20] J. Spencer, Countable Sparse Random Graphs, Random Structures and Algorithms *1* (1990), 205-214

[21] J. Spencer, Zero-One Laws via the Ehrenfeucht Game, Discrete Appl. Math. *30* (1991) 235-252

[22] J. Spencer, Sparse Random Graphs: A Continuum of Complete Theories *in* Proceedings of the International Conference "Sets, Graphs and Numbers", Colloq. Math. Soc. János Bolyai 60, Budapest 1991, D. Miklos, ed. , North Holland, publ. , pp 679-690

[23] J. Spencer, Zero-One Laws with Variable Probability, Journal of Symbolic Logic *58* (1993), 1-14

[24] J. Spencer and K. StJohn, Random Unary Predicates: Almost Sure Theories and Countable Models, Random Structures & Algorithms *13* (1998), 229-248

[25] J. Spencer and G. Tardos, Ups and Downs of First Order Sentences on Random Graphs, Combinatorica (submitted)

[26] P. Winkler, Random structures and zero-one laws, *Finite and Infinite Combinatorics in Sets and Logic*, N.W. Sauer, R.E. Woodrow and B. Sands, eds., NATO Advanced Science Institutes Series, Kluwer Academic Publishers, Dordrecht (1993), 399–420.