such potential H' there a.s. remain $\Theta(n^{v-\alpha e})$, hence at least one, x with $cl_a(x) \cong H$.

Now, in general, consider the $(r+1)$-st move. We set $b = a_{r+1}$, $a = a_r$ for notational convenience and recall $a = K + b$ where K is an upper bound on $cl_b(z_1, \ldots, z_{r+1})$. Points $x_1, \ldots, x_r \in G_1$, $y_1, \ldots, y_r \in G_2$ have been selected with

$$cl_a(x_1, \ldots, x_r) \cong cl_a(y_1, \ldots, y_r)$$

Spoiler picks, say, $x_{r+1} \in G_1$. We distinguish two cases. We say Spoiler has moved Inside if

$$x_{r+1} \in cl_K(x_1, \ldots, x_r)$$

Otherwise we say Spoiler has moved Outside.

Suppose Spoiler moves Inside. Then

$$cl_b(x_1, \ldots, x_r, x_{r+1}) \subseteq cl_{K+b}(x_1, \ldots, x_r) = cl_a(x_1, \ldots, x_r)$$

The isomorphism from $cl_a(x_1, \ldots, x_r)$ to $cl_a(y_1, \ldots, y_r)$ sends x_{r+1} to some y_{r+1} which Duplicator selects.

Suppose Spoiler moves Outside. Set $H = cl_b(x_1, \ldots, x_r, x_{r+1})$. Let H_0 be the union of all rigid extensions of any size of $x_1, \ldots, x_r$ in H. If $x_{r+1} \in H_0$ then, as $|H| \leq K$, $x_{r+1} \in cl_K(x_1, \ldots, x_r)$ and Spoiler moved Inside. Hence $x_{r+1} \notin H_0$. Since $|H| \leq K \leq a$, H_0 lies inside $cl_a(x_1, \ldots, x_r)$. The isomorphism between $cl_a(x_1, \ldots, x_r)$ and $cl_a(y_1, \ldots, y_r)$ maps H_0 into a copy of itself in the graph G_2.

For any copy of H_0 in G_2, let $N(H_0)$ denote the number of extensions of H_0 to H. From Theorem 3.2 one can show that a.s all $N(H_0) = \Theta(n^{v-\alpha e})$, with $v = v(H_0, H)$, $e = e(H_0, H)$ and $v - \alpha e > 0$. For a given H_0 each y_{r+1} is in only a bounded number of copies of H since all copies of H lie in $cl_b(y_1, \ldots, y_r, y_{r+1})$. Hence there are $\Theta(n^{v-\alpha e})$ vertices y_{r+1} so that $cl_b(y_1, \ldots, y_r, y_{r+1})$ contains H. Arguing as with the first move there a.s. are $\Theta(n^{v-\alpha e})$, hence at least one, y_{r+1} with $cl_b(y_1, \ldots, y_r, y_{r+1}) \cong H$. Duplicator selects such a y_{r+1}.

close α may be approximated by rationals of denominator at most t. This is often the case. If, for example, $\frac{1}{2}+\frac{1}{s+1}<\alpha<\frac{1}{2}+\frac{1}{s}$ then a.s. there will be two points $x_1, x_2 \in G(n, n^{-\alpha})$ having s common neighbors so that $|cl_1(x_1, x_2)| = s+2$.

Now we define the $a_1, \ldots, a_t$ of the lookahead strategy by reverse induction. We set $a_t = 0$. If at the end of the game Duplicator can assure that the 0-types of $x_1, \ldots, x_t$ and $y_1, \ldots, y_t$ are the same then they have the same induced subgraphs and he has won. Suppose, inductively, that $b = a_{r+1}$ has been defined. Let, applying the Lemma, K be a.s. an upper bound on all $cl_b(z_1, \ldots, z_{r+1})$. We then define $a = a_r$ by $a = K + b$.

Now we need show that a.s. this strategy works. Let $G_1 \sim G(n, n^{-\alpha})$, $G_2 \sim G(m, m^{-\alpha})$ and suppose Duplicator tries to play the $(a_1, \ldots, a_t)$ lookahead strategy on $EHR(G_1, G_2, t)$.

Set $a = a_1$ and consider the first move. Spoiler will select, say, $y = y_1 \in G_2$. Duplicator then must play $x = x_1 \in G_1$ with $cl_a(x) \cong cl_a(y)$. Can he always do so - that is, do a.s. G_1 and G_2 have the same values of $cl_a(x)$? The size of $cl_a(x)$ is a.s. bounded so it suffices to show for any potential H that either there almost surely is an x with $cl_a(x) \cong H$ or there almost surely is no x with $cl_a(x) \cong H$.

Let H have v vertices and e edges. Suppose H has a subgraph H' (possibly H itself) with v' vertices, e' edges and $v' - \alpha e' < 0$. The expected number of copies of H' in G_1 is

$$\Theta(n^{v'}p^{e'}) = \Theta(n^{v'-\alpha e'}) = o(1)$$

so a.s. G_1 contains no copy of H', hence no copy of H, hence no x with $cl_a(x) \cong H$. If this does not occur then (since, critically, α is irrational) all $v' - \alpha e' > 0$ so the expected number of copies of all such H' approaches infinity. From Theorem 1.4.5 a.s. G_1 has $\Theta(n^{v-\alpha e})$ copies of H. For x in appropriate position in such a copy of H we cannot deduce $cl_a(x) \cong H$ but only that $cl_a(x)$ contains H as a subgraph. (Essentially, x may have additional extension properties.) For each such x as $cl_a(x)$ is bounded, $cl_a(x)$ contains only a bounded number of copies of H. Hence there are $\Theta(n^{v-\alpha e})$ *different* $x \in G_1$ so that $cl_a(x)$ contains H as a subgraph.

Let H' be a possible value for $cl_a(x)$ that contains H as a subgraph. Let H' have v' vertices and e' edges. As (x, H') is rigid, (H, H') is dense and so

$$(v' - v) - \alpha(e' - e) < 0$$

There are $\Theta(n^{v'-\alpha e'})$ different x with $cl_a(x)$ containingH' but since $v' - \alpha e' < v - \alpha e$ this is $o(n^{v-\alpha e})$. Subtracting off such x for all the boundedly many

Example. Taking $\alpha \sim .51$ (but irrational, of course), $cl_1(x_1, x_2)$ consists of x_1, x_2 and all y adjacent to both of them. $cl_3(x_1, x_2)$ has those points and all y_1, y_2, y_3 which together with x_1 form a K_4 (note that this gives an (R, H) with $v = 3, e = 6$) and a finite number of other possibilities.

We can already describe the nature of Duplicator's strategy. At the end of the r-th move, with $x_1, \ldots, x_r$ and $y_1, \ldots, y_r$ having been selected from the two graphs, Duplicator will assure that these sets have the same $a_r - type$. We shall call this the $(a_1, \ldots, a_t)$ *lookahead strategy.* Here a_r must depend only on t, the total number of moves in the game and α. We shall set $a_t = 0$ so that at the end of the game, if Duplicator can stick to the $(a_1, \ldots, a_t)$ lookahead strategy then he has won. If, however, Spoiler picks, say, x_r so that there is no corresponding y_r with $x_1, \ldots, x_r$ and $y_1, \ldots, y_r$ having the same a_r-type then the strategy fails and we say that Spoiler wins. The values a_r give the "lookahead" that Duplicator uses but before defining them we need some preliminary results.

Lemma 4.6 Let α, $r, t > 0$ be fixed. Then there exists $K = K(\alpha, r, t)$ so that in $G(n, n^{-\alpha})$ a.s.

$$|cl_t(x_1, \ldots, x_r)| \leq K$$

for *all* $x_1, \ldots, x_r \in G$.

Proof. Set $K = r + t(L-1)$. If $X = \{x_1, \ldots, x_r\}$ has t-closure with more than K points then there will be L sets $Y^1, \ldots, Y^L$ disjoint from X, all $|Y^j| \leq t$ so that each $(X, X \cup Y^j)$ forms a rigid extension and with each Y^j having at least one point not in $Y^1 \cup \ldots Y^{j-1}$. Begin with X and add the Y^j in order. Adding Y^j will add, say, v_j vertices and e_j edges. Since $(X, X \cup Y^j)$ was *rigid*, $(X \cup Y^1 \cup \ldots \cup Y^{j-1}, X \cup Y^1 \cup \ldots \cup Y^j)$ is dense and so $v_j - e_j\alpha < 0$. As $v_j \leq t$ there are only a finite number of possible values of $v_j - e_j\alpha$ and so there is an $\epsilon = \epsilon(\alpha, r, t)$ so that all $v_j - e_j\alpha \leq -\epsilon$. Pick L (and therefore K) so that $r - L\epsilon < 0$. The existence of a t-closure of size greater than K would imply the existence in $G(n, n^{-\alpha})$ of one of a finite number of graphs that would have some $r + v_1 + \ldots + v_L$ vertices and at least $e_1 + \ldots + e_L$ edges. But the probability of G containing such a graph is bounded by

$$\begin{array}{ll} n^{r+v_1+\ldots+v_L} p^{e_1+\ldots+e_L} & = n^{r+v_1+\ldots+v_L-\alpha(e_1+\ldots+e_L)} \\ n^{r+(v_1-\alpha e_1)+\ldots+(v_L-\alpha e_L)} & \leq n^{r-L\epsilon} \\ = o(1) & \end{array}$$

so a.s. no such t-closures exist. □

Remark. The value of K given by the above proof depends strongly on how the disjunction over all distinct $u_1, \ldots, u_a, v_1, \ldots, v_b \in G$ with $a + b \leq s$. There are less than $s^2 n^s$ such choices as we can choose a, b and then the vertices. Thus

$$\Pr[E] \leq s^2 n^s (1-\epsilon)^{n-s}$$

But

$$\lim_{n\to\infty} s^2 n^s (1-\epsilon)^{n-s} = 0$$

and so E holds almost never. Thus $\neg E$, which is precisely the statement that $G(n,p)$ has the full level s extension property, holds almost always. □

But now we have proven Theorem 4.1. For any $p \in (0,1)$ and any fixed s as $m, n \to \infty$ with probability approaching one both $G(n,p)$ and $H(m,p)$ will have the full level s extension property and so Duplicator will win $EHR[G(n,p), H(m,p), s]$.

Why can't Duplicator use this strategy when $p = n^{-\alpha}$? We illustrate the difficulty with a simple example. Let $.5 < \alpha < 1$ and let Spoiler and Duplicator play a three move game on G, H. Spoiler thinks of a point $z \in G$ but doesn't tell Duplicator about it. Instead he picks $x_1, x_2 \in G$, both adjacent to z. Duplicator simply picks $y_1, y_2 \in H$, either adjacent or not adjacent dependent on whether $x_1 \sim x_2$. But now wily Spoiler picks $x_3 = z$. $H \sim H(m, m^{-\alpha})$ does not have the full level 2 extension property. In particular, most pairs y_1, y_2 do not have a common neighbor. Unless Duplicator was lucky, or shrewd, he then cannot find $y_3 \sim y_1, y_2$ and so he loses. This example does not say that Duplicator will lose with perfect play - indeed, we will show that he almost always wins with perfect play - it only indicates that the strategy used need be more complex. Now let us fix $\alpha \in (0,1)$, α irrational.

Now recall our notion of rooted graphs (R, H) but this time from the perspective of a particular $p = n^{-\alpha}$. We say (R, H) is *dense* if $v - e\alpha < 0$ and *sparse* if $v - e\alpha > 0$. The irrationality of α assures us that all (R, H) are in one of these categories. We call (R, H) *rigid* if for all S with $R \subseteq S \subset V(H)$, (S, H) is dense.

For any r, t there is a finite list (up to isomorphism) of rigid rooted graphs (R, H) containing r roots and with $v(R, H) \leq t$. In any graph G we define the t-closure $cl_t(x_1, \ldots, x_r)$ to be the union of all $y_1, \ldots, y_v$ with (crucially) $v \leq t$ which form an (R, H) extension where (R, H) is rigid. If there are no such sets we define the default value $cl_t(x_1, \ldots, x_r) = \{x_1, \ldots, x_r\}$. We say two sets $x_1, \ldots, x_r$ and $x'_1, \ldots, x'_r$ have the same t-type if their t-closures are isomorphic. (To be precise, these are ordered r-tuples and the isomorphism must send x_i into x'_i.)

since when the Zero-one Law is not satisfied $\lim_{n\to\infty} \Pr[G(n,p(n)) \models A]$ might not exist. If there is a subsequence n_i on which the limit is $c \in (0,1)$ we may use the same argument. Otherwise there will be two subsequences n_i, m_i on which the limit is zero and one respectively. Then letting $n, m \to \infty$ through n_i, m_i respectively, Spoiler will win $EHR[G,H,t]$ with probability approaching one. $\square$

Theorem 4.4 provides a bridge from Logic to Random Graphs. To prove that $p = p(n)$ satisfies the Zero-One Law we now no longer need to know anything about Logic - we just have to find a good strategy for the Duplicator.

We say that a graph G has the full level s extension property if for every distinct $u_1, \ldots, u_a, v_1, \ldots, v_b \in G$ with $a + b \leq s$ there is an $x \in V(G)$ with $\{x, u_i\} \in E(G)$, $1 \leq i \leq a$ and $\{x, v_j\} \notin V(G)$, $1 \leq j \leq b$. Suppose that G, H both have the full level $s-1$ extension property. Then Duplicator wins $EHR[G,H,s]$ by the following simple strategy. On the i-th round, with $x_1, \ldots, x_{i-1}, y_1, \ldots, y_{i-1}$ already selected, and Spoiler picking, say, x_i, Duplicator simply picks y_i having the same adjacencies to the $y_j, j < i$ as x_i has to the $x_j, j < i$. The full extension property says that such a y_i will surely exist.

Theorem 4.5 For any fixed p, $0 < p < 1$, and any s, $G(n,p)$ almost always has the full level s extension property.

Proof. For every distinct $u_1, \ldots, u_a, v_1, \ldots, v_b, x \in G$ with $a + b \leq s$ let $E_{u_1,\ldots,u_a,v_1,\ldots,v_b,x}$ be the event that $\{x, u_i\} \in E(G)$, $1 \leq i \leq a$ and $\{x, v_j\} \notin V(G)$, $1 \leq j \leq b$. Then

$$\Pr[E_{u_1,\ldots,u_a,v_1,\ldots,v_b,x}] = p^a(1-p)^b$$

Now define

$$E_{u_1,\ldots,u_a,v_1,\ldots,v_b} = \wedge_x \overline{E_{u_1,\ldots,u_a,v_1,\ldots,v_b,x}}$$

the conjunction over $x \neq u_1, \ldots, u_a, v_1, \ldots, v_b$. But these events are mutually independent over x since they involve different edges. Thus

$$\Pr[\wedge_x \overline{E_{u_1,\ldots,u_a,v_1,\ldots,v_b,x}}] = [1 - p^a(1-p)^b]^{n-a-b}$$

Set $\epsilon = \min(p, 1-p)^s$ so that

$$\Pr[\wedge_x \overline{E_{u_1,\ldots,u_a,v_1,\ldots,v_b,x}}] \leq (1-\epsilon)^{n-s}$$

The key here is that ϵ is a fixed (dependent on p, s) positive number. Set

$$E = \vee E_{u_1,\ldots,u_a,v_1,\ldots,v_b}$$

has t rounds. Each round has two parts. First the Spoiler selects either a vertex $x \in V(G)$ or a vertex $y \in V(H)$. He chooses which graph to select the vertex from. Then the Duplicator must select a vertex in the other graph. At the end of the t rounds t vertices have been selected from each graph. Let $x_1, \ldots, x_t$ be the vertices selected from $V(G)$ and $y_1, \ldots, y_t$ be the vertices selected from $V(H)$ where x_i, y_i are the vertices selected in the i-th round. Then Duplicator wins if and only if the induced graphs on the selected vertices are order-isomorphic: i.e., if for all $1 \leq i < j \leq t$

$$\{x_i, x_j\} \in E(G) \longleftrightarrow \{y_i, y_j\} \in E(H)$$

As there are no hidden moves and no draws one of the players must have a winning strategy and we will say that that player wins $EHR[G, H, t]$.
Lemma 4.3 For every First Order A there is a $t = t(A)$ so that if G, H are any graphs with $G \models A$ and $H \models \neg A$ then Spoiler wins $EHR[G, H, t]$.

A detailed proof would require a formal analysis of the First Order language so we give only an example. Let A be the property $\forall_x \exists_y [x \sim y]$ of not containing an isolated point and set $t = 2$. Spoiler begins by selecting an isolated point $y_1 \in V(H)$ which he can do as $H \models \neg A$. Duplicator must pick $x_1 \in V(G)$. As $G \models A$, x_1 is not isolated so Spoiler may pick $x_2 \in V(G)$ with $x_1 \sim x_2$ and now Duplicator cannot pick a "duplicating" y_2.
Theorem 4.4 A function $p = p(n)$ satisfies the Zero-One Law if and only if for every t, letting $G(n, p(n)), H(m, p(m))$ be independently chosen random graphs on disjoint vertex sets

$$\lim_{m,n \to \infty} \Pr[\text{ Duplicator wins} EHR[G(n, p(n)), H(m, p(m)), t]] = 1$$

Remark. For any given choice of G, H somebody must win $EHR[G, H, t]$. (That is, there is no random play, the play is perfect.) Given this probability distribution over (G, H) there will be a probability that $EHR[G, H, t]$ will be a win for Duplicator, and this must approach one.
Proof. We prove only the "if" part. Suppose $p = p(n)$ did not satisfy the Zero-One Law. Let A satisfy

$$\lim_{n \to \infty} \Pr[G(n, p(n)) \models A] = c$$

with $0 < c < 1$. Let $t = t(A)$ be as given by the Lemma. With limiting probability $2c(1-c) > 0$ exactly one of $G(n, p(n)), H(n, p(n))$ would satisfy A and thus Spoiler would win, contradicting the assumption. This is not a full proof and having radius at most two

$$\exists_x \forall_y [\neg(y = x) \wedge \neg(y \sim x) \longrightarrow \exists_z [z \sim y \wedge y \sim x]]$$

For any property A and any n, p we consider the probability that the random graph $G(n,p)$ satisfies A, denoted

$$\Pr[G(n,p) \models A]$$

Our objects in this section will be the theorem of Glebskii et.al. [1969] and independently Fagin[1976]
Theorem 4.1 For any fixed p, $0 < p < 1$ and any First Order A

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

and that of Shelah and Spencer[1988]
Theorem 4.2 For any *irrational* α, $0 < \alpha < 1$, setting $p = p(n) = n^{-\alpha}$

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

Both proofs are only outlined.

We shall say that a function $p = p(n)$ satisfies the Zero-One Law if the above equality holds for every First Order A.

The Glebskii/Fagin Theorem has a natural interpretation when $p = .5$ as then $G(n,p)$ gives equal weight to every (labelled) graph. It then says that any First Order property A holds for either almost all graphs or for almost no graphs. The Shelah/Spencer Theorem may be interpreted in terms of threshold functions. For example, $p = n^{-2/3}$ is a threshold function for containment of a K_4. That is, when $p \ll n^{-2/3}$, $G(n,p)$ almost surely does not contain a K_4 whereas when $p \gg n^{-2/3}$ it almost surely does contain a K_4. In between, say at $p = n^{-2/3}$, the probability is between 0 and 1, in this case $1 - e^{-1/24}$. The (admittedly rough) notion is that *at* a threshold function the Zero-One Law will not hold and so to say that $p(n)$ satisfies the Zero-One Law is to say that $p(n)$ is not a threshold function - that it is a boring place in the evolution of the random graph, at least through the spectacles of the First Order language. In stark terms: What happens in the evolution of $G(n,p)$ at $p = n^{-\pi/7}$? The answer: Nothing!

Our approach to Zero-One Laws will be through a variant of the Ehrenfeucht Game, which we now define. Let G, H be two vertex disjoint graphs and t a positive integer. We define a perfect information game, denoted $EHR[G,H,t]$, with two players, denoted Spoiler and Duplicator. The game family of extensions. Thus with probability $1 - o(n^{-1})$ there is a maximal disjoint family of extensions F with $|s-\mu| < \epsilon\mu$. As F consists of extensions

$$\Pr[N(x) < (1-\epsilon)\mu] = o(n^{-1})$$

To complete the upper bound we need show that $N(x)$ will not be much larger than $|F|$. Here we use only that $p = n^{-2/3+o(1)}$. There is $o(n^{-1})$ probability that $G(n,p)$ has an edge $\{x, x'\}$ lying in ten triangles. There is a $o(n^{-1})$ that $G(n,p)$ has a vertex x with $u_i, v_i, w_i, 1 \le i \le 7$ all distinct and all x, u_i, v_i and x, v_i, w_i triangles. When these do not occur $N(x) \le |F| + 70$ for any maximal disjoint family of extensions $|F|$ and so for any $\epsilon' > \epsilon$

$$\Pr[N(x) > (1+\epsilon')\mu] < o(n^{-1}) + \Pr[\text{some } |F| > (1+\epsilon)\mu] = o(n^{-1})$$

With some additional work one can find K so that the conclusions of the theorem hold for any $p = p(n)$ with $\mu > K \log n$. The general result is stated in terms of rooted graphs. For a given rooted graph (R, H) let $N(x_1, \ldots, x_r)$ denote the number of $(y_1, \ldots, y_v)$ giving an (R, H) extension. Set $\mu = \binom{n-r}{v} p^e$, the expected value of N in $G(n,p)$.
Theorem 3.2. Let (R, H) be strictly balanced. Then for all $\epsilon > 0$ there exists K so that if $p = p(n)$ is such that $\mu > K \log n$ then almost surely

$$|N(x_1, \ldots, x_r) - \mu| < \epsilon\mu$$

for all $x_1, \ldots, x_r$.

In particular if $\mu \gg \log n$ then almost surely all $N(x_1, \ldots, x_r) \sim \mu$.

4 Zero-One Laws

In this section we restrict our attention to graph theoretic properties expressible in the First Order theory of graphs. The language of this theory consists of variables $(x, y, z, \ldots)$, which always represent vertices of a graph, equality and adjacency $(x = y, x \sim y)$, the usual Boolean connectives $(\wedge, \neg, \ldots)$ and universal and existential quanfication $(\forall_x, \exists_y)$. Sentences must be finite. As examples, one can express the property of containing a triangle

$$\exists_x \exists_y \exists_z [x \sim y \wedge x \sim z \wedge y \sim z]$$

having no isolated point

$$\forall_x \exists_y [x \sim y]$$

3 All Vertices in nearly the same number of Triangles

Returning to the example of §1, let $N(x)$ denote the number of triangles containing vertex x. Set $\mu = \binom{n-1}{2}p^3$ as before.
Theorem 3.1. For every $\epsilon > 0$ there exists K so that if $p = p(n)$ is such that $\mu = K \log n$ then almost surely

$$(1-\epsilon)\mu < N(x) < (1+\epsilon)\mu$$

for *all* vertices x.

We shall actually show that for a given vertex x

$$\Pr[|N(x) - \mu| > \epsilon\mu] = o(n^{-1})$$

If the distribution of $N(x)$ were Poisson with mean μ then this would follow by Large Deviation results and indeed our approach will show that $N(x)$ is closely approximated by the Poisson distribution.

We call F a maximal disjoint family of extensions if F consists of pairs $\{x_i, y_i\}$ such that all x, x_i, y_i are triangles in $G(n,p)$, the x_i, y_i are all distinct, and there is no $\{x', y'\}$ with x, x', y' a triangle and x', y' both distinct from all the x_i, y_i. Let $Z^{(s)}$ denote the number of maximal disjoint families of size s. Lets restrict $0 \le s \le \log^2 n$ (a technical convenience) and bound $E[Z^{(s)}]$. There are $\sim \binom{n-1}{2}^s / s!$ choices for F. Each has probability $(p^3)^s$ that all x_i, y_i do indeed give extensions. We further need that the $n - 1 - 2s \sim n$ other vertices contain no extension. The calculation of §1 may be carried out here to show that this probability is $\sim e^{-\mu}$. All together

$$E[Z^{(s)}] \le (1 + o(1))\frac{\mu^s}{s!}e^{-\mu}$$

But now the right hand side is asymptotically the Poisson distribution so that we can choose K so that

$$\sum * E[Z^{(s)}] = o(n^{-1}) \qquad (*)$$

where $\sum^*$ is over $s < \log^2 n$ with $|s - \mu| > \epsilon\mu$.

When $s > \log^2 n$ we ignore the condition that F be maximal so that $E[Z^{(s)}] < \mu^s/s! = o(n^{-10})$, say. Thus $(*)$ holds with $\sum^*$ over all s with $|s-\mu| > \epsilon\mu$. Thus with probability $1 - o(n^{-1})$ all maximal disjoint families of extensions F have $|s - \mu| < \epsilon\mu$. But there *must* be some maximal disjoint

2 Rooted Graphs

The above result was only a special case of a general result of Spencer[1990] which we now state. By a *rooted graph* is meant a pair (R, H) consisting of a graph $H = (V(H), E(H))$ and a specified proper subset $R \subset V(H)$ of vertices called the roots. For convenience let the vertices of H be labelled $a_1, \ldots, a_r, b_1, \ldots, b_v$ with $R = \{a_1, \ldots, a_r\}$. In a graph G we say that vertices $y_1, \ldots, y_v$ make an (R, H)-extension of vertices $x_1 \ldots, x_r$ if all these vertices are distinct; y_i, y_j are adjacent in G whenever b_i, b_j are adjacent in H; and x_i, y_j are adjacent in G whenever a_i, b_j are adjacent in H. So G on $x_1, \ldots, x_r, y_1, \ldots, y_v$ gives a copy of H which may have additional edges – except that edges between the x's are not examined. We let $Ext(R, H)$ be the property the for all $x_1, \ldots, x_r$ there exist $y_1, \ldots, y_v$ giving an (R, H) extension. For example, when H is a triangle and R one vertex $Ext(R, H)$ is the statement that every vertex lies in a triangle. When H is a path of length t and R the endpoints $Ext(R, H)$ is the statement that every pair of vertices lie on a path of length t. When $R = \emptyset$ $Ext(\emptyset, H)$ is the already examined statement that there exists a copy of H. As in that situation we have a notion of balanced and strictly balanced. We say (R, H) has type (v, e) where v is the number of nonroot vertices and e is the number of edges of H, not counting edges with both vertices in R. For every S with $R \subset S \subseteq V(H)$ let (v_S, e_S) be the type of $(R, H|_S)$. We call (R, H) balanced if $e_S/v_S \leq e/v$ for all such S and we call (R, H) strictly balanced if $e_S/v_S < e/v$ for all proper $S \subset V(H)$. We call (R, H) nontrivial if every root is adjacent to at least one nonroot.
Theorem 2.1. Let (R, H) be a nontrivial strictly balanced rooted graph with type (v, e) and $r = |R|$. Let c_1 be the number of graph automorphism $\sigma : V(H) \rightarrow V(H)$ with $\sigma(x) = x$ for all roots x. Let c_2 be the number of bijections $\sigma : R \rightarrow R$ which are extendable to some graph automorphism $\lambda : V(H) \rightarrow V(H)$. Let $\mu > 0$ be arbitrary and fixed. Let $p = p(n)$ satisfy

$$\frac{n^v p^e}{c_1} = \ln\left(\frac{n^r}{c_2 \mu}\right)$$

Then

$$\lim_{n\rightarrow\infty} \Pr[G(n,p) \models Ext(R, H)] = e^{-\mu}$$

While the counting of automorphisms leads to some technical complexities the proof is essentially that of the "every vertex in a triangle" case.

the number of vertices x not lying in a triangle. Then from Linearity of Expectation

$$E[X] = \sum_{x \in V(G)} E[X_x] \rightarrow c$$

We need show that the Poisson Paradigm applies to X. To do this we show that all moments of X are the same as for the Poisson distribution. Fix r. Then

$$E[X^{(r)}/r!] = S^{(r)} = \sum \Pr[C_{x_1} \wedge \ldots \wedge C_{x_r}],$$

the sum over all sets of vertices $\{x_1, \ldots, x_r\}$. All r-sets look alike so

$$E[X^{(r)}/r!] = \binom{n}{r} \Pr[C_{x_1} \wedge \ldots \wedge C_{x_r}] \sim \frac{n^r}{r!} \Pr[C_{x_1} \wedge \ldots \wedge C_{x_r}]$$

where $x_1, \ldots, x_r$ are some particular vertices. But

$$C_{x_1} \wedge \ldots \wedge C_{x_r} = \wedge \overline{B_{x_i yz}},$$

the conjunction over $1 \leq i \leq r$ and all y, z. We apply Janson's Inequality to this conjunction. Again $\epsilon = p^3 = o(1)$. The number of $\{x_i, y, z\}$ is $r\binom{n-1}{2} - O(n)$, the overcount coming from those triangles containing two (or three) of the x_i. (Here it is crucial that r is fixed.) Thus

$$\sum \Pr[B_{x_i yz}] = p^3 \left(r\binom{n-1}{2} - O(n) \right) = r\mu + O(n^{-1+o(1)})$$

As before Δ is p^5 times the number of pairs $x_i yz \sim x_j y' z'$. There are $O(rn^3) = O(n^3)$ terms with $i = j$ and $O(r^2 n^2) = O(n^2)$ terms with $i \neq j$ so again $\Delta = o(1)$. Therefore

$$\Pr[C_{x_1} \wedge \ldots \wedge C_{x_r}] \sim e^{-r\mu}$$

and

$$E[X^{(r)}/r!] \sim \frac{(ne^{-\mu})^r}{r!} = \frac{c^r}{r!}$$

Hence X has limiting Poisson distribution, in particular $\Pr[X = 0] \rightarrow e^{-\mu}$.
$\square$

Lecture 5: Counting Extensions and Zero-One Laws

The threshold behavior for the existence of a copy of H in $G(n,p)$ is well understood. Now we turn to what, in a logical sense, is the next level which we call *extension statements*. We want $G(n,p)$ to have the property that every $x_1, \ldots, x_r$ belong to a copy of H. For example ($r = 1$), every vertex lies in a triangle. We find the fine threshold behavior for this property and further show - continuing this example - that for p a bit larger almost surely every vertex lies in about the same number of triangles.

1 Every Vertex in a Triangle

Let A be the property that every vertex lies in a triangle.
Theorem 1.1. Let $c > 0$ be fixed and let $p = p(n)$, $\mu = \mu(n)$ satisfy

$$\binom{n-1}{2} p^3 = \mu$$

$$e^{-\mu} = \frac{c}{n}$$

Then

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

Proof. First fix $x \in V(G)$. For each unordered $y, z \in V(G) - \{x\}$ let B_{xyz} be the event that $\{x, y, z\}$ is a triangle of G. Let C_x be the event $\wedge \overline{B_{xyz}}$ and X_x the corresponding indicator random variable. We use Janson's Inequality to bound $E[X_x] = \Pr[C_x]$. Here $p = o(1)$ so $\epsilon = o(1)$. $\sum Pr[B_{xyz}] = \mu$ as defined above. Dependency $xyz \sim xuv$ occurs if and only if the sets overlap (other than in x). Hence

$$\Delta = \sum_{y,z,z'} \Pr[B_{xyz} \wedge B_{xyz'}] = O(n^3)p^5 = o(1)$$

since $p = n^{-2/3+o(1)}$. Thus

$$E[X_x] \sim e^{-\mu} = \frac{c}{n}$$

Now define

$$X = \sum_{x \in V(G)} X_x,$$