

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:

The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

MEASURES ON THE RANDOM GRAPH

by

Michael H. Albert

Department of Mathematics
Carnegie Mellon University
Pittsburgh, PA 15213

Research Report No. 92-144 ₂

March, 1992

510.6
C28R
92-144

Measures on the random graph

Michael H. Albert
Department of Mathematics
Carnegie Mellon University
Pittsburgh, PA 15213

March 18, 1992

Abstract

We consider the problem of characterizing the finitely additive probability measures on the definable subsets of the random graph which are invariant under the action of the automorphism group of this graph. We show that such measures are all integrals of Bernoulli measures (which arise from the coin-flipping model of the construction of the random graph). We also discuss generalizations to other theories.

1 Introduction

Let Ω denote the countable random graph, a countable model of T^{ind} , or the countable homogeneous universal graph – these are all the same thing expressed in different dialects of mathematics (according to whether your “kitchen culture” comes from [1], [8], or [2].)

A simple concrete description (from [2]) is to take the vertex set of Ω to be

$$\omega = \{0, 1, 2, \dots\},$$

and to have x and y (with $x < y$) adjacent if and only if 2^z occurs in the unique representation of y as a sum of powers of 2. Because Ω is an \aleph_0 -categorical structure, its automorphism group has only finitely many orbits on n -tuples for any finite n . More precisely, if the induced subgraphs of Ω on two sequences of vertices:

$$\begin{aligned} \mathbf{a} &= a_1, a_2, \dots, a_n; \\ \mathbf{b} &= b_1, b_2, \dots, b_n \end{aligned}$$

are isomorphic (with the isomorphism sending each a_i to b_i for $1 \leq i \leq n$), then there exists $g \in \text{Aut}(\Omega)$ such that $a_i g = b_i$ for $1 \leq i \leq n$, or more succinctly:

$$\mathbf{a}g = \mathbf{b}.$$

Note that we will write functions on the right.

A subset A of Ω is called *definable* if for some sequence \mathbf{b} of vertices, and some first order formula $\phi(x, \mathbf{y})$ in the language of a single binary relation E (whose interpretation is the edge relation in the random graph),

$$A = \{a \in \Omega : \Omega \models \phi(a, \mathbf{b})\}.$$

The field of definable subsets of Ω will be denote $\text{Def}(\Omega)$. This field is generated by the finite sets, and the sets defined by formulas $\phi_{n,k}$ for $0 \leq k \leq n$ of the following form:

$$\phi_{n,k}(x, y_1, y_2, \dots, y_n) := \bigwedge_{j=1}^k xEy_j \wedge \bigwedge_{j=k+1}^n \neg xEy_j.$$

For a sequence of parameters \mathbf{b} , the set defined by $\phi_{n,k}(x, \mathbf{b})$ which we denote $\phi_{n,k}(\Omega, \mathbf{b})$ is just the set of points which are adjacent to b_1, b_2, \dots, b_k and independent of b_{k+1}, \dots, b_n . In fact, an axiomatization of Ω is given by the set of axioms:

$$\forall \mathbf{y} \exists x \phi_{n,k}(x, \mathbf{y})$$

for each $n \geq 0$ and each $0 \leq k \leq n$.

For a general discussion of \aleph_0 -categorical structures, their automorphism groups, and the structure of their definable sets (in particular the quantifier elimination which yields the results above) we suggest [2].

Another method of constructing Ω is to begin with a single vertex, and to form a sequence of finite graphs, adding a vertex at a time, and deciding which of the previous vertices to connect it to by flipping a fair coin for each edge. The union of the resulting graphs will then be isomorphic to Ω with probability 1. From this model, we see that there is a finitely additive probability measure $\mu_{1/2}$ on the field of definable subsets of Ω determined by:

$$\mu_{1/2}(\phi_{n,k}(\Omega, \mathbf{b})) = 1/2^n.$$

Moreover this measure has the property that it is invariant under the action of the automorphism group of Ω namely, if A and B are definable sets and $A = Bg$ for some $g \in \text{Aut}(\Omega)$ then $\mu_{1/2}(A) = \mu_{1/2}(B)$. For the "random" viewpoint of the random graph, see [1].

The goal of this paper is to classify all finitely additive measures on Ω which have this property (invariance under the automorphism group of Ω .) In [2] another

type of measure on Ω is considered. The connection between such measures and the ones which we deal with, as well as measures on other types of structures are considered in Section 4.

2 Definitions

We have defined the random graph Ω and its field of definable subsets $Def(\Omega)$ above. In this paper we consider functions:

$$\mu : Def(\Omega) \rightarrow \mathbf{R}$$

with the following properties:

- For all A in $Def(\Omega)$, $\mu(A) \geq 0$.
- If A and B are disjoint elements of $Def(\Omega)$ then $\mu(A \cup B) = \mu(A) + \mu(B)$.
- If A is an element of $Def(\Omega)$ and g is an automorphism of Ω then $\mu(Ag) = \mu(A)$.
- $\mu(\Omega) = 1$.

We refer to such functions as finitely additive automorphism invariant probability measures on Ω , or more briefly simply as measures on Ω .

For any finite subset F of Ω and any measure μ on Ω , $\mu(F) = 0$ since F has infinitely many disjoint images under the automorphism group of Ω .

If G is a graph, then $V(G)$ denotes the set of vertices of G ,

Fix now a measure μ on Ω . Let G be a finite graph, and let X be a subset of the vertices of G . Then we can find a sequence $\mathbf{b} \in \Omega$ such that there is an isomorphism between the induced subgraph on \mathbf{b} and G which sends an initial segment of \mathbf{b} to X (this is by the universality of the random graph). Moreover, if \mathbf{b} and \mathbf{c} are two such sequences, then there is an automorphism g of Ω with $\mathbf{b} = c g$ (by homogeneity). Hence the quantity:

$$\mu(G, X) := \mu(\phi_{n,k}(\Omega, \mathbf{b})) \quad \text{where } n = |G| \text{ and } k = |X|$$

is well defined. Furthermore if α is an isomorphism from G to another graph H then:

$$\mu(G, X) = \mu(H, X\alpha) \tag{1}$$

since G and H may be viewed as subgraphs of Ω and α can be extended to an automorphism of Ω . Clearly, the quantities $\mu(G, X)$ as G runs over finite graphs (or over a representative set of isomorphism types of finite graphs) determine all the quantities $\mu(\phi_{n,k}(\Omega, \mathbf{b}))$.

Henceforth we will work only with the quantities $\mu(G, X)$. Also given any finite graph, we implicitly identify it with some isomorphic subgraph of Ω .

Let G be any finite graph and let X be any subset of $V(G)$. Suppose that G' is a graph which extends G by a single new vertex v . By the additivity of μ :

$$\mu(G, X) = \mu(G', X) + \mu(G', X \cup \{v\}). \quad (2)$$

This relation just states that any vertex of Ω is either adjacent to v or independent of v , and that these possibilities are disjoint. An extension of this relation which applies to any graph G' containing G is the following:

$$\mu(G, X) = \sum_{Y \subseteq V(G') - V(G)} \mu(G', X \cup Y). \quad (3)$$

Of course this follows easily from (2) and induction.

It is also easy to see that:

$$\begin{aligned} \sum_{X \subseteq G} \mu(G, X) &= 1 \\ \mu(G, X) &\geq 0 \end{aligned} \quad (4)$$

Any function $\mu(G, X)$ which satisfies the relations (1), (2), and (4) is determined by a unique measure on Ω . So in our attempt to classify measures on Ω we will concentrate on classifying the functions with these properties. We will prove:

Theorem 1 *Let a measure μ on Ω be given. Then there is a unique probability measure ν on $[0, 1]$ such that for all finite graphs G and all $X \subseteq V(G)$:*

$$\mu(G, X) = \int_0^1 p^{|X|} (1-p)^{|G|-|X|} d\nu.$$

In particular, $\mu(G, X)$ depends only on $|G|$ and $|X|$ and not on the structure of G .

For the measure $\mu_{1/2}$ considered above, the associated probability measure is an atomic measure concentrated at $1/2$. More generally, for $0 \leq p \leq 1$ there are measures μ_p which arise (at least for $p \neq 0, 1$) from a construction of the random graph with an unfair coin, and which correspond to atomic measures concentrated at p . The content of theorem 1 is that every measure defined on Ω lies in the closure of the convex hull of the set

$$\{\mu_p : p \in [0, 1]\}.$$

In effect this means that “the random graph can only be constructed by flipping coins”.

In the proof of this theorem it will be necessary to add many vertices to certain finite graphs G . To this end we define for each positive integer N , the set

$$[N] = \{1, 2, \dots, N\}$$

and we assume that $[N]$ and $V(G)$ are disjoint whenever necessary (since all the concepts we deal with are isomorphism-invariant this is not a restrictive assumption.) Also K_N denotes the complete graph with vertex set $[N]$.

The next section is devoted to the proof of Theorem 1, and the following section to some further discussion and generalizations.

3 Proof of Theorem 1

The proof of Theorem 1 is rather involved. For the reader's guidance we begin with a plan of the proof:

- There is a unique probability measure ν^K such that for all n

$$\mu(K_n, X) = \int_0^1 p^{|X|} (1-p)^{n-|X|} d\nu^K,$$

where K_n is the complete graph on n vertices. (Lemma 2)

- If:

$$\nu^K = \sum_{i=1}^n \lambda_i \nu_i$$

where the λ_i are non-negative, and the ν_i are probability measures which are supported on intervals with at most one point in common then

$$\mu = \sum_{i=1}^n \lambda_i \mu_i$$

where the μ_i are $Aut(\Omega)$ invariant probability measures and:

$$\mu_i^K = \nu_i.$$

This is Lemma 3, the decomposition lemma. Some extra but inessential technical assumptions are required.

- If G is fixed, and the support of ν^K is sufficiently narrow, then for all $X \subseteq V(G)$:

$$\mu(G, X) \approx \int_0^1 p^{|X|} (1-p)^{|G|-|X|} d\nu.$$

This is Lemma 4. Again there is a small technical restriction which is dealt with in Lemma 5.

- Hence by decomposing ν^K into narrow pieces, and taking limits:

$$\mu(G, X) = \int_0^1 p^{|X|} (1-p)^{|G|-|X|} d\nu.$$

as claimed in Theorem 1.

The first part of the proof shows that the measure behaves properly when restricted to complete graphs. The main work then comes in extending this behavior to arbitrary graphs. In both the second and the third part of the proof, the idea is to begin with an arbitrary finite graph G , and to add many more vertices so that the resulting graph G^* contains a large clique, and has automorphisms which permute the members of the clique arbitrarily. We can then apply equation (3) to calculate $\mu(G, X)$ from some quantities $\mu(G^*, Y)$. The presence of the large clique in G^* permits us to find bounds for some of these quantities by relating them to $\mu(K_N, Z)$ (Since $K_N \subseteq G^*$ this is possible by further application of equation (3)). From there it is just a matter of applying some fairly easy limiting arguments to get the results we desire.

There are a number of places in the argument where N tends to ∞ and δ to 0, or where we deduce that some quantity has an appropriate limit by showing that a number of other quantities behave reasonably and yield a suitable approximation. To handle these arguments in full ϵ - δ detail would make the proofs more unreadable than they already are, so we often give such arguments quite informally.

Lemma 2 *Let μ be a measure on Ω . Then there is a unique probability measure ν^K on $[0, 1]$ such that:*

$$\mu(K_n, X) = \int_0^1 p^{|X|} (1-p)^{n-|X|} d\nu^K,$$

for all complete graphs K_n and subsets X of their vertices.

Proof: By the isomorphism-invariance of μ , $\mu(K_n, X)$ can depend only on n and $k = |X|$. Let $\mu(n, k)$ denote this value, and define

$$\nu(l, k) = \mu(l+k, k).$$

Then from equation (2) (applied to the passage from K_n to K_{n+1}) and non-negativity we get:

$$\begin{aligned} \nu(0, 0) &= 1 \\ \nu(l, k) &\geq 0 \quad \text{for all } 0 \leq l, k \\ \nu(l, k) &= \nu(l+1, k) + \nu(l, k+1). \end{aligned}$$

But:

$$0 \leq \nu(1, k) = \nu(0, k) - \nu(0, k + 1),$$

and subsequent iterates of this establish that for all $r \geq 0$:

$$(-\Delta)^r(\nu(0, 0), \nu(0, 1), \dots, \nu(0, k), \dots) \geq 0$$

where Δ is the difference operator:

$$\Delta(x_0, x_1, x_2, \dots, x_n, \dots) = (x_1 - x_0, x_2 - x_1, \dots, x_{n+1} - x_n, \dots).$$

But by Theorem 2 in Section 7.3 of [6] (the Hausdorff moment theorem) this implies that there is a unique probability measure ν^K on $[0, 1]$ such that:

$$\nu(k, l) = \int_0^1 x^k (1-x)^l d\nu^K$$

or,

$$\mu(n, k) = \int_0^1 x^k (1-x)^{n-k} d\nu$$

as claimed. ■

Our second lemma will allow us to decompose μ given a decomposition of ν^K .

To decompose μ on G we first expand G to a graph G_N by adding a large complete graph connected to all vertices of G . Then we associate each term which arises in the computation of $\mu(G, X)$ in the expanded graph G_N to a unique part of the ν^K decomposition. Next we define $\mu_i(G, X)$ just by taking the sum of the terms which arise from ν_i . Finally a certain amount of house cleaning is required to check that the properties required of a measure still hold and we're done.

We fix a decomposition of ν^K of the following type:

$$\nu^K = \lambda_1 \nu_1 + \lambda_2 \nu_2 + \dots + \lambda_n \nu_n$$

where each ν_i is a probability measure on $[0, 1]$, and there is a sequence:

$$0 = a_0 < a_1 < a_2 < \dots < a_n = 1$$

such that the support of ν_i is contained in $[a_{i-1}, a_i]$, and none of the points a_i for $0 < i < n$ are atoms of ν . For simplicity we will also assume that the a_i for $0 < i < n$ are all irrational.

Given a positive integer N we define $I_i(N)$ to be the collection of subsets Y of $[N]$ for which

$$a_{i-1}N \leq |Y| \leq a_i N,$$

and for $N' > N$ define:

$$P_i(N, N') = \{Y \subseteq [N'] : |Y| \in I_i(N') \text{ and } |Y \cap [N]| \in I_i(N)\}$$

An important fact which we will use is that for N a sufficiently large positive integer any $N' > N$, and any i between 1 and n :

$$\int_0^1 \left(\sum_{Y \in P_i(N, N')} p^{|Y|} (1-p)^{N'-|Y|} \right) d\nu = \lambda_i + o(1) \quad (5)$$

(where the $o(1)$ is with respect to $N \rightarrow \infty$.) This is true because the polynomial being integrated converges to the characteristic function of $[a_{i-1}, a_i]$ as N tends to infinity. Note that

$$\int_0^1 p^{|Y|} (1-p)^{N'-|Y|} d\nu = \mu(K_{N'}, Y).$$

Note also that it follows that any sum of distinct terms $\mu(K_{N'}, Y)$ where either $Y \cap [N] \in I_i(N)$ or $Y \in I_i(N')$ but Y is not an element of $P_i(N, N')$ is $o(1)$ since the sum in equation (5) over $Y \subseteq I_i(N)$ or $Y \subseteq I_i(N')$ is also $\lambda_i + o(1)$.

Now we can state and prove the lemma:

Lemma 3 *Let μ be a measure on Ω and suppose that the associated probability measure ν^K is decomposed as above. Then there are measures μ_i such that for all G and X :*

$$\mu(G, X) = \sum_{i=1}^n \lambda_i \mu_i(G, X)$$

and $\mu_i^K = \nu_i$.

Proof: Fix for the moment a graph G , and define G_N to be the graph obtained by adding vertices $[N]$ to G which form a clique, and are connected to every vertex in G . For $X \subseteq G$ define:

$$\mu_i^N(G, X) = \sum_{Y \in I_i(N)} \mu(G_N, X \cup Y)$$

We claim that the limit as N tends to infinity of $\mu_i^N(G, X)$ exists. If $N' > N$ then:

$$\mu_i^N(G, X) = \sum_Y \sum_{Z \subseteq [N'] - [N]} \mu(G_N, X \cup Y \cup Z)$$

where the first sum is over the same index set as before. After cancellation of common terms, the difference between this and $\mu_i^{N'}(G, X)$ is the difference between:

$$\sum_W \mu(G'_N, X \cup W)$$

where $W \cap [N] \in I_i(N)$ but $W \notin I_i(N')$,

$$\sum_V \mu(G'_N, X \cup V).$$

where $V \cap [N] \notin I_i(N)$, but $V \in I_i(N')$. Each of these is dominated by the corresponding sum over $K_{N'}$ of terms of the form:

$$\mu(K_{N'}, W) \text{ or } \mu(K_{N'}, V).$$

But both of these sums are $o(1)$ by the note after equation (5). Now define

$$\mu_i(G, X) = (1/\lambda_i) \lim_{N \rightarrow \infty} \mu_i^N(G, X).$$

Clearly

$$\mu(G, X) = \sum_{i=1}^n \lambda_i \mu_i(G, X)$$

so it remains only to show that each μ_i is a measure. The only non-trivial part to check is equation (2), but this follows immediately from the fact that if G' extends G by a single vertex v , then:

$$(G')_N = (G_N)'$$

so the terms which occur in the sums for $\mu_i^N(G', X)$ and $\mu_i^N(G', X \cup v)$, are exactly the same as those obtained by applying equation (2) to the terms in $\mu_i^N(G, X)$. ■

Now we will prove that if ν^K is narrow, then for a fixed graph G ,

$$\mu(G, X) \approx \int_0^1 p^{|X|} (1-p)^{|G|-|X|} d\nu.$$

It suffices to show that:

$$\mu(G, X) \approx \bar{p}^{|X|} (1-\bar{p})^{|G|-|X|}$$

where

$$\int_0^1 p d\nu^K = \bar{p}.$$

Define the *discrepancy* of μ on G to be:

$$Dis_G(\mu) = \sup_{Y \subseteq G} |\mu(G, Y) - \bar{p}^{|Y|} (1-\bar{p})^{|G|-|Y|}|,$$

and define the width of a probability measure on $[0, 1]$ to be the length of the smallest closed interval which contains the support of the measure.

Lemma 4 *Let G be a finite graph. Let $\epsilon > 0$ be given. There exists a $\delta > 0$ such that if μ is a measure on Ω , and the width of ν^K is less than δ then $\text{Dis}_G(\mu) < \epsilon$ (i.e. the discrepancy tends to 0 as the width tends to 0).*

Proof: If $\bar{p} = 0$ or $\bar{p} = 1$ then ν^K is a point mass at 0 or 1 respectively, and it is easy to check that in these extreme cases μ is essentially trivial (if $\bar{p} = 0$ then $\mu(G, X) = 0$ whenever X is non-empty, and if $\bar{p} = 1$ then $\mu(G, X) = 0$ whenever $X \neq G$.) In these cases the result is true, and so we assume that $p \neq 0, 1$ and set

$$\bar{q} = \bar{p}/(1 - \bar{p}).$$

Let v be any vertex of G , and let N be a positive integer. Construct the graph G_N^v by adding new vertices $[N]$ to G which, together with v form a clique, and which are connected to another vertex $w \neq v$ of G if and only if v is. Note that any permutation of the vertices $\{v\} \cup [N]$ extends to an automorphism of G (the remaining vertices will be fixed.)

Let X be any subset of G which does not contain v . Then:

$$\mu(G, X) = \sum_{Y \subseteq [N]} \mu(G_N^v, X \cup Y) = \sum_{j=0}^N \binom{N}{j} \mu(G_N^v, X \cup [j]). \quad (6)$$

Where the last equality is true because for any j element subset Y of $[N]$ there is an automorphism of G_N^v which fixes G and sends Y to $[j]$. Similarly:

$$\mu(G, X \cup \{v\}) = \sum_{j=0}^N \binom{N}{j} \mu(G_N^v, X \cup \{v\} \cup [j]). \quad (7)$$

Now for every j between 0 and N

$$\mu(G_N^v, X \cup [j]) \leq \mu(K_N, [j]),$$

and

$$\mu(G_N^v, X \cup \{v\} \cup [j]) \leq \mu(K_N, [j])$$

because both the left hand sides of the inequality are terms which arise in the computation of the right hand side, when K_N is expanded by adding a copy of G to form G_N^v .

In particular, since the support of ν^K is contained in $[\bar{p} - \delta, \bar{p} + \delta]$ we may conclude that:

$$\sum_{j=0}^N \binom{N}{j} \mu(G_N^v, X \cup [j]) = \sum_{j \in [N(\bar{p}-\delta), N(\bar{p}+\delta)]} \binom{N}{j} \mu(G_N^v, X \cup [j]) + o(1),$$

and

$$\sum_{j=0}^N \binom{N}{j} \mu(G_N^v, X \cup \{v\} \cup [j]) = \sum_{j \in [N(\bar{p}-\delta), N(\bar{p}+\delta)]} \binom{N}{j} \mu(G_N^v, X \cup [j]) + o(1),$$

and if $0 < \bar{p} - \delta < \bar{p} + \delta < 1$ then the extreme terms of either sum are also $o(1)$. The following argument applies only to this case.

If the discrepancy of G were 0 then we would have

$$\bar{q} \mu(G, X) - \mu(G, X \cup \{v\}) = 0.$$

So we aim to show that the quantity on the left hand side is small.

$$\begin{aligned} |\bar{q} \mu(G, X) - \mu(G, X \cup \{v\})| &= \\ \sum_{j \in [N(\bar{p}-\delta), N(\bar{p}+\delta)]} \left(\bar{q} \binom{N}{j} - \binom{N}{j-1} \right) \mu(G_N^v, X \cup [j]) &+ o(1). \end{aligned}$$

We may choose f_δ which is $o(1)$ as $\delta \rightarrow 0$ so that for sufficiently large N and all $j \in [N(\bar{p} - \delta), N(\bar{p} + \delta)]$,

$$\left| \bar{q} \binom{N}{j} - \binom{N}{j-1} \right| \leq f_\delta \binom{N}{j}.$$

So:

$$|\bar{q} \mu(G, X) - \mu(G, X \cup v)| \leq f_\delta + o(1) = g_\delta$$

Where

$$g_\delta = o(1) \text{ as } \delta \rightarrow 0$$

can be chosen independently of v and X .

Applying the triangle inequality we can conclude that for any $Y \subseteq V(G)$:

$$|\bar{q}^{|Y|} \mu(G, \emptyset) - \mu(G, Y)| \leq \left(\sum_{k=0}^{|Y|-1} \bar{q}^k \right) g_\delta \leq h_\delta$$

where h_δ is also $o(1)$ as $\delta \rightarrow 0$ since the sum is bounded by a number depending only on \bar{q} and $|G|$. But now by summing all the quantities $\mu(G, Y)$ we will be able to show that $\mu(G, \emptyset)$ must be close to $(1 - \bar{p})^{|G|}$, and then that $\mu(G, Y)$ must be close to $\bar{p}^{|Y|} (1 - \bar{p})^{|G| - |Y|}$ provided that δ is sufficiently small, and thus the discrepancy of G will be small for sufficiently small δ .

What about the case when $\bar{p} - \delta < 0$? If ν^K has no atom at 0 then the above argument can easily be modified to work in this case (since we can effectively ignore any contribution from terms corresponding to sufficiently small subsets.)

However, the case where ν^K does have an atom at 0 (or at 1) must be dealt with separately, and this is done in the next lemma. ■

We need to eliminate the annoying special case in the previous lemma. This is easily done:

Lemma 5 *Suppose that ν^K has an atom of weight λ at 0. Then*

$$\mu = \lambda\mu_1 + (1 - \lambda)\mu_2$$

where μ_1 and μ_2 are measures and $\mu_1(G, X) = 0$ unless $X = \emptyset$.

Proof: It is easy to see that it suffices to show that

$$\mu(G, \emptyset) \geq \lambda$$

for all G . Let G_N be the graph obtained by joining adding a clique of size N to G , each vertex of which is adjacent to every vertex of G . Since

$$\mu(G, \emptyset) \geq \mu(G_N, \emptyset)$$

it suffices to show that we can make the latter arbitrarily close to λ for sufficiently large N . But:

$$\mu(K_N, \emptyset) = \mu(G_N, \emptyset) + \sum_{X \subseteq G, X \neq \emptyset} \mu(G_N, X).$$

Each term in the sum is dominated by $\mu(K_{N+1}, \{x\})$ and since:

$$\mu(K_N, \emptyset) = \mu(K_{N+1}, \emptyset) + \mu(K_{N+1}, \{x\})$$

we can make such a term arbitrarily small. Since the number of such terms depends only on $|G|$ and not on N we are finished. ■

A similar argument deals with atoms at 1. So in order to prove Theorem 1 it is sufficient to consider the case where ν^K does not have an atom at 0 or 1, so that Lemma 4 can be applied to any term of a decomposition of ν^K . Actually the proof of the theorem is now essentially complete.

Let μ , and G be given. We want to show that:

$$\mu(G, X) = \int_0^1 p^{|X|}(1-p)^{|G|-|X|} d\nu^K.$$

And of course it suffices to show that $\mu(G, X)$ is arbitrarily close to the right hand side. Since ν^K can have only countably many atoms, we can decompose

$$\nu^K = \sum_{i=1}^n \lambda_i \nu_i$$

so that each ν_i is as narrow as we like, and the conditions of the decomposition lemma are also met (since we assumed that ν^K has no atom at 0 or 1)

$$\mu^K = \sum_{i=1}^n \lambda_i \mu_i$$

and $\mu_i(G, X)$ is close to (say within ϵ of)

$$\int_0^1 p^{|X|} (1-p)^{|G|-|X|} d\nu_i$$

for each $1 \leq i \leq n$. But then

$$\begin{aligned} & \left| \mu(G, X) - \int_0^1 p^{|X|} (1-p)^{|G|-|X|} d\nu \right| \leq \\ & \sum_{i=1}^n \lambda_i \left| \mu_i(G, X) - \int_0^1 p^{|X|} (1-p)^{|G|-|X|} d\nu_i \right| \\ & \leq \epsilon. \end{aligned}$$

So we have finished.

4 Discussion and other examples

On page 112 of [2], P. Cameron introduces a different type of measure p on Ω . We will call his measures “subgraph measures” because for a given finite graph G , $p(G)$ is intended to indicate the probability that a sequence \mathbf{a} in Ω is isomorphic (via a specified labelling) to G . With this interpretation, he then gives the following natural axioms for such a function p :

$$\begin{aligned} p(A) & \geq 0 \text{ for all } A; \\ p(\emptyset) & = 1; \\ p(A) & = \sum p(A') \text{ where the sum is over all children } A' \text{ of } A; \\ p(A) & \text{ is independent of the labelling of } A. \end{aligned}$$

Note that a child of A is an extension of A by a single vertex.

These conditions are superficially quite similar to those required of a measure μ . However the complicated recurrence for subgraph measures as opposed to the much simpler one:

$$\mu(G, X) = \mu(G', X) + \mu(G', X \cup \{v\})$$

means that an analysis of all subgraph measures leading to a theorem akin to the one above seems out of the question (Cameron himself remarks that the conditions are too general and admit “too many” solutions).

From the logical point of view, the subgraph measures are really a sequence of measures on the n -types over the empty set, satisfying the conditions that the measure of a type should be the sum of the measures of all its extensions, and an invariance under permutation of variables. This suggests that such measures could be investigated more fruitfully in theories where the collection of n -types is much simpler than is the case for graphs.

There is an obvious way to attempt to define a subgraph measure from a measure, namely by the use of conditional probabilities. If G' is an extension of G by a vertex v which is adjacent to X then we should have:

$$\mu(G, X) = \frac{p(G')}{p(G)} \quad \text{or} \quad p(G') = p(G)\mu(G, X).$$

However, an arbitrary p satisfying the conditions above will not in general yield a measure μ , nor will an arbitrary μ yield a suitable p . Considering the second construction only, problems arise because of the requirement of isomorphism invariance. In the end $p(G')$ will be a product of factors of the form $\mu(G_i, X)$, and the order in which the vertices are chosen must not affect the value of this product. From this idea it is not hard to prove:

Proposition 6 *If a subgraph measure p and a measure μ on Ω are connected by the relationship:*

$$\mu(G, X) = \frac{p(G')}{p(G)} \quad \text{or} \quad p(G') = p(G)\mu(G, X).$$

Then there exists $\bar{p} \in [0, 1]$ such that:

$$\mu(G, X) = \bar{p}^{|X|}(1 - \bar{p})^{|G| - |X|}.$$

In other words both p and μ arise from the model of Ω obtained by flipping a fixed coin with a probability \bar{p} of “heads”.

Proof: Let μ and p be given. Let I_k denote the k element independent set. Let ν be the probability measure on $[0, 1]$ associated to μ and define:

$$\bar{p}_n = \int_0^1 p^n d\nu.$$

Note that we may assume that $p_1 \neq 0, 1$ for the result is trivial in those cases.

We wish to show that $\bar{p}_n = \bar{p}^n$ for every n . The result holds for $n = 0, 1$. Assume that it holds for all $n < N$. Consider the graph S_{N+1} which is a star

on $N + 1$ vertices (i.e. an independent set of size N and another, central, vertex connected to all of these.) By adding the central vertex last we see that:

$$p(S_{N+1}) = p(I_{N-1}) \left(\int_0^1 (1-p)^{N-1} d\nu \right) \left(\int_0^1 p^N d\nu \right).$$

Adding the central vertex second last we get:

$$p(S_{N+1}) = p(I_{N-1}) \left(\int_0^1 p^{N-1} d\nu \right) \left(\int_0^1 p(1-p)^{N-1} d\nu \right).$$

Equating the two results (and noting $p(I_{N-1}) \neq 0$) and using the inductive hypothesis we get:

$$(1-\bar{p})^{N-1} \bar{p}_N = \bar{p}^{N-1} (\bar{p}(1-\bar{p})^{N-1} - (-1)^{N-1} p^N + (-1)^{N-1} p_N).$$

Hence

$$p_N ((1-\bar{p})^{N-1} - (-1)^{N-1} \bar{p}^{N-1}) = p^N ((1-\bar{p})^{N-1} - (-1)^{N-1} \bar{p}^{N-1})$$

and so $p_N = \bar{p}^N$ as required. ■

The question of classification of finitely additive probability measures on other \aleph_0 -categorical structures, or of finding a general theory of such, is currently under investigation. In many cases, it is easy or trivial to classify the measures. For example for the theory of dense linear order without endpoints, a measure is determined by a single numerical parameter $\lambda \in [0, 1]$ which is the measure of any (and hence all) sets of the form $\{x : x > a\}$.

A somewhat more interesting case to consider is that of triangle-free graphs. If one considers the first order sentences which almost all finite triangle free graphs satisfy then the countable model of this theory is an infinite bipartite graph. In this graph the two parts are both infinite sets, and given any two disjoint subsets of a single part there exists a vertex in the other part adjacent to all the vertices in the first set and none in the second. This graph is not homogeneous since there are independent sets which contain vertices from both parts, and independent sets which are contained in a single part, and there can be no isomorphism of the graph which exchanges such sets. This is the only obstruction to homogeneity; in fact if two subgraphs are isomorphic as bipartite graphs with specified bipartition, then there is an automorphism extending the isomorphism. There is also an automorphism exchanging the parts. So for any measure, each part will have measure $1/2$, and there will be a unique probability measure ν such for disjoint finite subsets X and Y of one part:

$$\mu(X \cup Y, X) = \int_0^1 p^{|X|} (1-p)^{|Y|} d\nu.$$

The proof is very easy since any two finite subsets of the same size, each contained in a single part are conjugate. Thus one needs only a version of Lemma 2, which is proved in exactly the same way.

The result “almost all finite graphs are bipartite” simplified the example above. Another candidate graph to consider is the infinite homogeneous-universal triangle free graph, Ω_3 . This is a countable graph, unique up to isomorphism, with the following properties:

- It is triangle free;
- Given any finite subgraph G and any triangle free graph H containing G there is an embedding of H in Ω_3 which extends the identity map on G .

In Ω_3 the definable sets are also generated by finite sets, and those given by “extension conditions”. So a measure is determined by the values $\mu(G, X)$ as in Ω (where here we must add the guarantee that $\mu(G, X) = 0$ if two vertices in X are adjacent.)

It seems clear that on Ω_3 the values $\mu(G, X)$ must depend on the structure of G since for example $\mu(K_2, \{x\}) = 0$, while it would seem that the same need not be true if we replace K_2 by I_2 . On the other hand, the arguments for Ω seem also to be valid, though we must of course work with large independent sets rather than large cliques. As the following shows, the second viewpoint is correct, though for an unfortunate reason:

Proposition 7 *Let μ be a measure on Ω_3 . Then:*

$$\mu(G, X) = \begin{cases} 1 & \text{if } X = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

In particular $\mu(G, X)$ depends only on $|G|$ and $|X|$.

Proof: Certainly if G and X are non-empty:

$$\mu(G, X) \leq \mu(K_1, \{x\}).$$

Consider the graph, G_N , consisting of a single vertex $\{x\}$ adjacent to an independent set of size N . Then for any non-empty subset Y of the independent set,

$$\mu(G_N, \{x\} \cup Y) = 0.$$

Hence:

$$\mu(K_1, \{x\}) = \mu(G_N, \{x\}) \leq \mu(I_N, \emptyset).$$

Suppose that $\mu(I_N, \emptyset) \rightarrow 0$ as $N \rightarrow \infty$. Then $\mu(G, X) = 0$ for all non-empty X . But in that case $\mu(I_N, \emptyset) = 1$ for all N , a contradiction.

But the argument of Lemma 5 can be used to show that if $\mu(I_N, \emptyset) \rightarrow \lambda \neq 0$ as $N \rightarrow \infty$, then unless $\lambda = 1$, a new measure is obtained by subtracting λ from $\mu(G, \emptyset)$ (and renormalizing) for all G , and this new measure has the paradoxical property from above. So the only possibility is that $\mu(I_n, \emptyset) = 1$ for all n and hence (easily) that μ is as described above. ■

An alternative proof of this result is more directly connected to the abstract study of measures on models. First consider a sequence of pairs of vertices (a_i, b_i) in Ω_3 with the following property a_i is adjacent to b_j if and only if $i < j$. Now consider the sets S_i of common neighbors of a_i and b_i . Then for $i < j$, S_i and S_j are disjoint since a_i and b_j cannot have a common neighbor. On the other hand the measure of all the sets S_i must be the same, since each is just the set of neighbors of a pair of independent points. So this measure must be 0. But now let c_i be any sequence of independent points, and let N_i be the neighbors of c_i . By the above, the measure of $N_i \cap N_j$ is 0 when $i \neq j$, so the N_i are “almost disjoint” sets of equal measure, and hence must have measure 0. But then by monotonicity, the measure of any set which is contained in the set of neighbors of a point is 0, which gives the stated result.

Thus there are no non-trivial measures on Ω_3 . A similar argument also works for any other infinite universal-homogeneous graph (these are all like Ω_3 only they omit either a clique or an independent set of some fixed size n . This classification is given in [7]), though the almost sure models (see [5]) of these theories all behave nicely.

An unfortunate result of this argument is that one gains no new insight into the question of whether the theory of Ω_3 has the finite model property. A structure M has the finite model property if every sentence true of M is true of some finite structure. The random graph Ω has this property, as does the naked set (which is in some sense Ω_2). It is an open question whether or not Ω_k has the finite model property for $k \geq 3$ (some examples of graphs which satisfy a few of the required sentences are known: [3]). An interesting measure on Ω_3 would have provided some guidance as to where to look for examples. For example had it turned out that $1/5 < \mu(K_1, \{x\}) < 1/3$ then it would have been reasonable to look for examples among graphs with n vertices where the average degree was around $n/4$. But the failure of interesting measures to exist yields no such positive guidance. It suggests, but by no means proves that the edge density of finite graphs satisfying many of the extension axioms should be small.

Doug Ensley [4] has begun to establish a connection between the existence of interesting measures, and various conditions which arise in stability theory. An oversimplification of his results is “stable = dull, strict order property = dull or bad, independence property = maybe interesting”. The “maybe” is necessary as the example of Ω_3 shows.

5 Acknowledgements

Neil Calkin deserves a great deal of credit for this work. He was the partner in a number of unsuccessful attempts at a proof of theorem 1. Had he not departed for warmer climes he would most probably have been a co-author. The following submitted graciously to relentless interrogation by the author – who couldn't remember much of the basic measure theory needed for the proof of theorem 1: V. Mizel, R. Nicolaides, J. Schaeffer, M. Soner, L. Tartar, and N. Walkington.

References

- [1] Béla Bollobás, *Random Graphs*, Academic Press, London, 1985.
- [2] P. Cameron, *Oligomorphic Permutation Groups*, London Math. Soc. Lect. Note Ser., v. 152, Cambridge University Press, 1990.
- [3] G. Cherlin, *Three-extensible triangle free graphs and geometries*, private communication, 1991.
- [4] D. Ensley, Ph.D. thesis, Carnegie Mellon University (in preparation)
- [5] P. Erdős, D.J. Kleitman, and Rothschild B.L., Asymptotic enumeration of K_n -free graphs, In *Teoria Combinatorie* vol. II, Academia Naz. die Lincei, Rome, 1976, pp. 19–27.
- [6] W. Feller, *An introduction to probability theory and its applications, vol. II*, John Wiley and Sons, New York, 1966.
- [7] A. H. Lachlan, and R.E. Woodrow, Countable ultrahomogeneous undirected graphs, *Trans. Amer. Math. Soc.*, 284, (1980), 431–461
- [8] S. Shelah, *Classification theory and the number of non-isomorphic models*, Revised Edition, North-Holland, Amsterdam, 1990.

FEB 13 2004

Carnegie Mellon University Libraries



3 8482 01373 5713