Notes on Random Graphs: The purpose of these notes is to explain what is meant by Paul Erdos' result that "any two random graphs are isomorphic.". These notes are structured in such a way that we avoid talking about randomness and probability until the last section. Most of the proofs in these notes have nothing (directly) to do with random graphs. To save words, we stipulate now that all our infinite graphs have the natural numbers $\boldsymbol{N}$ as their vertex sets.

The Extension Property: Let Γ be an infinite graph. We say that Γ has the *extension property* if the following is true. For any two disjoint finite subsets $A, B \subset \boldsymbol{N}$, there exists a vertex $v \in \boldsymbol{N} - A - B$ such that v is connected to all vertices in A and to no vertices in B.

Lemma 0.1 *If Γ has the extension property then Γ also has the following property. For any disjoint finite subsets $A, B \subset \boldsymbol{N}$, there are infinitely many vertices $\{v_i\}$ so that v_i is connected to all vertices in A and none in B.*

Proof: Assume that we have produced $v_1, ..., v_n$ having the above property with respect to A and B. We then let $A' = A \cup v_1 \cup ... \cup v_n$ and $B' = B$. The extension property guarantees that there is some v' such that v' is connected to all vertices in A' and none in B'. But then we set $v_{n+1} = v'$, and $v_1, ..., v_{n+1}$ all have the desired property with respect to A and B. By induction, we can find an infinite list of such vertices. ♠

Just to anticipate where these notes are going, we're going to prove later on that a random graph with a countable vertex set has (with probability 1) the extension property. So, any results about graphs with the extension property will then hold (with probability 1) for random graphs.

Isomorphisms between Graphs Here is the main result.

Lemma 0.2 *Suppose that Γ_1 and Γ_2 are two infinite graphs, both having the extension property. Then Γ_1 and Γ_2 are isomorphic.*

Proof: Let $\Gamma_j(n)$ denote the subgraph induced by vertices $\{1, ..., n\}$ of Γ_j. Certainly there is an isomorphism $f_1 : \Gamma_1(1) \to \Gamma_2(1)$. We are just talking about single points here. We let $g_1 = f_1^{-1}$.

Suppose, by induction, that we have produced maps f_n and g_n and subgraphs $\Gamma_1'(n)$ and $\Gamma_2'(n)$ such that

- $\Gamma_1(n) \subset \Gamma'_1(n)$.
- $\Gamma_2(n) \subset \Gamma'_2(n)$.
- $f_n : \Gamma'_1(n) \to \Gamma'_2(n)$ is an isomorphism.
- $g_n : \Gamma'_2(n) \to \Gamma'_1(n)$ is the inverse of f_n.

If f_n is already defined on $(n+1)$, we let $(i_{n+1}) = f_n(n+1)$. Otherwise, we make the following definition. Let A_1 be the set of vertices of $\Gamma'_1(n)$ connected in Γ_1 to $(n+1)$. Let B_1 be the vertices of $\Gamma'_1(n)$ not connected in Γ_1 to $(n+1)$. By the extension property, there is some vertex (i_{n+1}) of Γ_2 which is connected in Γ_2 to all vertices of $f_n(A_1)$ and not connected in Γ_2 to any vertices of $f_n(A_2)$.

If g_n is already defined on $(n+1)$, we define $(j_{n+1}) = g_n(n+1)$. Otherwise, we define (j_{n+1}) just like we defined (i_{n+1}), but with the roles of Γ_1 and Γ_2 reversed. We let $\Gamma'_1(n+1)$ be the graph induced by $V(\Gamma'_1(n)) \cup (j_{n+1})$ and we let $\Gamma'_2(n+1)$ be the graph induced by $V(\Gamma'_2(n)) \cup (i_{n+1})$. We let $f_{n+1} = f_n$ on $\Gamma'_1(n)$ and then (if necessary)

$$f_{n+1}(n+1) = (i_{n+1}), \qquad f_{n+1}(j_{n+1}) = (n+1).$$

Likewise we define $g_{n+1} = g_n$ on $\Gamma'_2(n)$ and (if necessary)

$$g_{n+1}(n+1) = (j_{n+1}), \qquad g_{n+1}(i_{n+1}) = (n+1).$$

We then map in the edges according to where their endpoints go. This is the desired extension.

By induction we get an infinite sequence $\{f_n\}$ and $\{g_n\}$ of maps, and the sequences are compatible in the sense that $f_n = f_{n+1}$ and $g_n = g_{n+1}$ wherever both maps are defined. For any vertex v of Γ_1 we define $f(v) = f_n(v)$, where n is any value such that v is a vertex of $\Gamma'_1(n)$. The compatibility between these maps guarantees that f is well defined. Similarly we define the limiting map g. By construction f and g are inverses of each other, and f is an isomorphism from Γ_1 to Γ_2 (and g is an isomorphism from Γ_2 to Γ_1.) ♠

It is worth pointing out why we considered both f and g in the above construction. Were we to just consider f we might have just produced an isomorphism between Γ_1 and some subgraph of Γ_2. By considering both $\{f_n\}$ and $\{g_n\}$ at the same time, we guaranteed that $f : \Gamma_1 \to \Gamma_2$ was onto.

Homogeneity: Our argument gives us somewhat more than we have claimed so far. Call a subgraph of a graph *induced* if it is, in the classical sense, induced by some subset of the vertices. An induced subgraph involves all the edges of the graph which connect vertices within the vertex set of the subgraph.

Lemma 0.3 *Suppose that* Γ_1 *and* Γ_2 *are two infinite graphs with the extension property. Suppose that* $\Gamma_1' \subset \Gamma_1$ *and* $\Gamma_2' \subset \Gamma_2$ *are finite induced subgraphs. If there is an isomorphism* $f' : \Gamma_1' \to \Gamma_2'$ *then there is an isomorphism* $f : \Gamma_1 \to \Gamma_2$ *which extends* f'.

Proof: We set $f_0 = f'$ and $g_0 = f_0^{-1}$. We then repeat the argument given in the previous lemma, producing sequence of maps $\{f_n\}$ and $\{g_n\}$ which respectively extend f_0 and g_0. The limit map f is the desired isomorphism. ♠

A graph is called *vertex transitive* if there is a graph automorphism that takes any vertex to any other vertex. A graph is called *edge transitive* if there is a graph automorphism which takes any edge to any other edge.

Corollary 0.4 *A graph with the extension property is both vertex transitive and edge transitive.*

Proof: Let Γ be such a graph. Let v_1 and v_2 be arbitrary vertices of Γ. We set $\Gamma_1 = \Gamma_2 = \Gamma$, and $\Gamma_1' = \{v_1\}$ and $\Gamma_2' = \{v_2\}$. We then apply the above lemma. This proves vertex transitivity. The same kind of argument establishes edge transitivity. ♠

The Rado Graph: Now we know that there is at most one countable graph with the extension property, and that this graph (assuming it exists) is both vertex transitive and edge transitive. However, we have yet to give an example of a graph like this. The *Rado graph* is the canonical example.

Suppose that $i < j \in \boldsymbol{N}$ are two numbers. We join i to j by an edge if and only if the ith bit in the binary expansion of j is a 1. We count the bits from right to left. Thus (k) is connected to 2^{k-1} vertices in a row, then skips 2^{k-1} vertices, then is connected to 2^{k-1} vertices, and so on. To show that the Rado graph has the desired properties, suppose that A and B are finite disjoint subsets of $\boldsymbol{N}$. Just choose any number that has all 1s in the places corresponding to numbers of A and all 0s in the places corresponding to numbers of B. The vertex with this number does the job.

Measure Theory: Here is a very brief and superficial account of the large subject of measure theory. Let $S \subset [0,1]$ be a set. A *countable open cover* of S is a countable collection $\{I_j\}$ of open [1] intervals in $[0,1]$ such that $S \subset \bigcup I_j$. Given S, we define

$$m(S) = \inf_{\mathcal{C}} \sum_{i=1}^{\infty} |I_j|. \tag{1}$$

Here $|I_j|$ is the length of the interval I_j. The infimum is taken over the set $\mathcal{C}$ of all open covers of S. We call S a *null set* if $m(S) = 0$. The countable union of null sets is also null.

The function m is defined for any subset of $[0,1]$ but it only behaves well for certain subsets. A *σ-algebra* of subsets of $[0,1]$ is a collection $\mathcal{A}$ of subsets with the following properties.

- If S belongs to $\mathcal{A}$ then so does $[0,1] - S$.
- The countable union of subsets in $\mathcal{A}$ belongs to $\mathcal{A}$.
- The countable intersection of subsets in $\mathcal{A}$ belongs to $\mathcal{A}$.

The intersection of any collection of σ-algebras is again a σ-algebra. For this reason, it makes sense to speak of the smallest σ-algebra that contains all open sets of $[0,1]$. This σ-algebra is called the *Borel σ-algebra*. A *Borel set* is a member of the Borel σ-algebra. Intuitively, a Borel set is obtained by starting with open sets and taking countable unions, complements, and intersections countably many times.

A *Lebesgue set* is any set of the form $B \cup N$, where B is a Borel set and N is a null set. When S is a Lebesgue set, $m(S)$ is called the *Lebesgue measure* of S. The function m behaves nicely for Lebesgue sets. For instance, if $\{S_i\}$ is any countable collection of pairwise disjoint Lebesgue sets, we have

$$m(\bigcup S_i) = \sum m(S_i).$$

Likewise, $m([0,1] - S) + m(S) = 1$ for any Lebesgue set S.

[1] We count $[0,x)$ and $(x,1]$ as open intervals in $[0,1]$.

If S is a set which has a length in the traditional sense – e.g., S is a finite union of intervals – then $m(S)$ coincides with the length of S. What is nice about the Lebesgue measure is that it makes sense for sets which don't have length in the traditional sense. For instance, the set of irrationals in $[0,1]$ has measure 1 because the set of rationals (a countable set) has measure 0.

Random Graphs: In this section, "measure" always denotes "Lebesgue measure". Let $\mathcal{S}$ denote the set of all infinite binary sequences which have infinitely many 0s and infinitely many 1s. There is a map ϕ from $\mathcal{S}$ to $[0,1]$, which maps the binary sequence to the number in $[0,1]$ that has this sequence as its binary expansion. ϕ is injective, and the image of ϕ contains everything in $[0,1]$ except a countable set. Let's write $[0,1]' = \phi(\mathcal{S})$. Note that $[0,1]'$ has measure 1.

Call an infinite graph Γ *normal* if Γ has infinitely many edges, and the complement of Γ has infinitely many edges. Choose some enumeration of the pairs $i,j \in \boldsymbol{N}$ with $i<j$. Each binary sequence $s \in \mathcal{S}$ gives rise to a normal infinite graph: If (i,j) is the nth pair, we join (i) to (j) in our graph if and only if the nth bit of s is 1. This construction clearly produces every normal graph. In this way, we produce a bijection between the set of normal infinite graphs and $[0,1]'$. If we have some subset of the set of normal graphs, we can talk about its measure. By this, we mean the measure of the corresponding set in $[0,1]'$.

Intuitively, a sequence in $\mathcal{S}$ represents the result of flipping a fair coin infinitely often. So, we are flipping a coin for each edge to decide if that edge belongs to our graph. The identification of $\mathcal{S}$ with $[0,1]'$ does exactly the right thing in terms of measuring the probabilities of events. For instance, the set in $[0,1]'$ corresponding to those sequences having n specified bits has measure 2^{-n}.

To say that a normal graph has some property with probability 1 is to say that the set of normal graphs having that property has measure 1. Or, in other words, the set of normal graphs not having that property has measure 0. We don't lose anything by throwing out the set of non-normal infinite graphs, because there are only countably many of these. Determining the edges by a fair coin flip produces a normal graph with probability 1.

Lemma 0.5 *With probability* 1*, a normal infinite graph has the extension property.*

Proof: Let A and B be finite disjoint subsets of $\boldsymbol{N}$. For any $v \in \boldsymbol{N}-A-B$, let $P(v)$ denote the set of normal graphs for which v fails to satisfy the basic property with respect to A and B. That is, either v is not connected to some vertex in A or else v is connected to some vertex in B. The set $P(v)$ corresponds to the subset of $[0,1]'$ consisting of numbers which have some bit "incorrect" in one of the spots determined by A or B. Letting n be the cardinality of $A \cup B$, we see that the set $P(v)$ has measure $1-2^{-n}$. Intuitively, we are flipping a coin n times and the chances that everything goes right is 2^{-n}. So, the chances that something goes wrong is $1-2^{-n}$.

What happens with respect to one vertex is independent from what happens with respect to another. Hence, for distinct vertices $v_1, ..., v_k$, we have

$$m(P(v_1) \cap ... \cap P(v_k)) = (1-2^{-n})^k.$$

But then the set $S(A,B)$ of normal graphs for which *all* the vertices fail to have the extension property for the pair (A,B) has measure at most $(1-2^{-n})^k$ for all $k = 1, 2, 3, ...$ This is only possible if $S(A,B)$ has measure 0.

In other words, if we fix A and B, the set of normal graphs having the extension property for A and B has measure 1. But there are countably many pairs (A,B) of finite disjoint sets. The intersection of a countable collection of measure 1 sets again has measure 1. Hence, the set of normal graphs having the extension property has measure 1. ♠

If we produce two random graphs then, with probability 1, they both have the extension property. Hence, they are isomorphic to each other (and also to the Rado graph.) This is what is meant by the statement that any two random countable graphs are isomorphic.