EXTENSION AND THE 0-1 LAW ON THE COUNTABLE RANDOM GRAPH

ZERLINA YAU

Abstract. The countable random graph is characterized by the extension property, which allows for the construction of partial isomorphisms. This property can be used to demonstrate various other features of random graphs, including uniqueness and a high degree of regularity. We then use the extension property to prove the 0-1 law for random graphs, which connects the infinite random graph to its finite counterparts.

Contents

1. Introduction

A *graph* is a structure consisting of a set of vertices joined by edges. If we let the vertex set be countably infinite, and select edges independently with probability $\frac{1}{2}$, then we obtain the *countable random graph*. Despite its entirely probabilistic construction, the random graph is in fact unique up to isomorphism, as was proven by Erdős and Rényi in 1959. As a consequence, the random graph is highly stable under certain finite changes, such that any such modification produces a graph isomorphic to the original. This result relies on a defining property of random graphs, dubbed the *extension property*. This is what enables us to demonstrate the "sameness" of all infinite random graphs. If we take the extension property as a sequence of finite statements, we can additionally draw conclusions about *finite* random graphs – as it turns out, random graphs adhere to a 0-1 law. This is a theorem from first-order logic which states that sentences are either almost always true, or almost always false in finite structures. For graphs in particular, whether

Date: September 3, 2025.

a sentence is likely to be true on a finite graph corresponds to its truth value on the infinite graph, with the probability approaching 0 or 1 as the vertex set grows large. There are multiple ways to prove this result – our discussion will go over proofs using the compactness theorem and the Ehrenfeucht-Fraïssé game.

2. Basic Notions of Graphs

Before speaking of the Erdős-Rényi model in specific, we will go over some definitions for general graphs.

Definition 2.1. A *graph* is a structure $G(V, E)$ equipped with a vertex set V, an edge set E, and a relation $\sim$. The edge set E is given by some subset of all unordered pairs of vertices, $E \subseteq \{\{v, w\} \mid v, w \in V, v \neq w\}$.

For vertices $v, w \in V$, we define the adjacency relation $\sim$ as $v \sim w$ if $\{v, w\} \in E$. If $v \sim w$, we say that v and w are *neighbors*.

Definition 2.2. Let $G(V, E)$ and $H(W, F)$ be two graphs, and let $f : V \to W$. We say f is an *isomorphism* if f is a bijection, and f preserves adjacency, i.e. if $v, v' \in V$, then $v \sim v'$ if and only if $f(v) \sim f(v')$

Note that if $f : V \to W$ is an isomorphism, then its inverse $f^{-1} : W \to V$ exists and is also an isomorphism. Thus, we simply say that G and H are *isomorphic*, denoted as $G \cong H$.

3. Some Background on Probability

As its name suggests, the random graph is a probabilistically-constructed object. As such, Erdős and Rényi's theorem is proven using a probabilistic construction as well, which will require some knowledge of probability theory. In today's discussion, we consider probability theory as an extension of measure theory. For further reading, see chapters 2 and 3 of Richard Bass's *Real Analysis For Graduate Students* [1].

Definition 3.1. Let X be a set. A *σ-algebra* $\mathcal{A}$ is a collection of subsets of X satisfying the following properties:

(1) $\emptyset \in \mathcal{A}$ and $X \in \mathcal{A}$;
(2) For $A \subseteq X$, denote its complement $A^c := X \setminus A$. If $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$ (closure under complement);
(3) If $A_1, A_2, \cdots$ is a countable sequence of sets in $\mathcal{A}$, then $\bigcup_{i=1}^{\infty} A_i \in \mathcal{A}$ (closure under countable union).

The ordered pair $(X, \mathcal{A})$ is called a *measurable space*. A set $A \subseteq X$ is *measurable* if $A \in \mathcal{A}$.

Note that an immediate implication of the above definition is that a countable *intersection* of sets in $\mathcal{A}$ is also in $\mathcal{A}$, as $\bigcap_{i=1}^{\infty} A_i = (\bigcup_{i=1}^{\infty} A_i^c)^c$.

Definition 3.2. Let X be a set and $\mathcal{A}$ a σ-algebra over X. A *measure* on $(X, \mathcal{A})$ is a function $\mu : \mathcal{A} \to [0, \infty]$ satisfying the the following properties:

(1) $\mu(\emptyset) = 0$;
(2) If $A_1, A_2, \cdots$ is a countable sequence of pairwise disjoint sets in $\mathcal{A}$, where for $i \neq j$, $A_i \cap A_j = \emptyset$, then $\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i)$ (countable additivity).

The ordered triple $(X, \mathcal{A}, \mu)$ is called a *measure space*.

The following two facts will be helpful for later.

Fact 3.3 (Monotonicity)**.** Let $(X, \mathcal{A}, \mu)$ be a measure space. If $A_1, A_2 \in \mathcal{A}$ with $A_2 \subseteq A_1$, then $\mu(A_2) \leq \mu(A_1)$.

Fact 3.4 (Union Bound)**.** Let $(X, \mathcal{A}, \mu)$ be a measure space and let $A_1, A_2, \cdots$ be measurable sets. Then

$$\mathbb{P}[\bigcup_{i=1}^{\infty} A_i] \leq \sum_{i=1}^{\infty} \mathbb{P}[A_i].$$

Fact 3.3 results from noticing that $A_1 = A_2 \cup (A_1 \setminus A_2)$, so by countable additivity and non-negativity of a measure,

$$\mu(A_1) = \mu(A_2) + \mu(A_1 \setminus A_2) \geq \mu(A_2).$$

Proof of Fact 3.4 can be found in chapter 3 of [1].

As it turns out, the notion of measure lends itself rather nicely to defining probability. We consider a sample space Ω, which we can think of as the set of all final outcomes. Specifically, we look at certain subsets of the sample space, which we call *events*. In particular, we want to be able to speak about cases when an event does *not* occur (closure under complement), or when combinations of events occur (closure under countable union). We can then assign events probabilities between 0 and 1 using a probability function $\mathbb{P}$ which satisfies countable additivity. Importantly, we specify $\mathbb{P}[\Omega] = 1$ so that the probabilities of an event and its complement add up to 1.

Definition 3.5. A *probability space* is a measure space $(\Omega, \mathcal{S}, \mathbb{P})$ such that $\mathbb{P}[\Omega] = 1$. If $(\Omega, \mathcal{S}, \mathbb{P})$ is a probability space and $S \subseteq \Omega$ with $S \in \mathcal{S}$, then we call S an *event*. (Note that $\mathbb{P}[S]$ denotes the (probabilistic) measure of event S.)

Example 3.6. Suppose we want to model the outcomes of a 6-sided die roll. We can consider the sample space of all possible faces, $\Omega = \{1, 2, 3, 4, 5, 6\}$, and want to assign probabilities such that each face has a probability of $\frac{1}{6}$. Then, our σ-algebra $\mathcal{S}$ is the set of all subsets of Ω, and the probability function $\mathbb{P}$ is given by $\mathbb{P}[S] = \frac{|S|}{6}$.

We'll also want to speak about *event independence*. This typically refers to events which are explicitly uncorrelated, such as the outcomes of two dice rolls. In order to calculate the probabilities of two or more independent events, we simply multiply the probabilities together. In probability theory, we don't always have a notion of whether events "should" be correlated, so independence is characterized solely by this multiplication rule. Thus, we have the following definition:

Definition 3.7. Let $S, T \subseteq \Omega$ be events. We say that S and T are *independent* if

$$\mathbb{P}[S \cap T] = \mathbb{P}[S]\mathbb{P}[T].$$

Example 3.8. Consider the probability space of two independent coin flips, where the sample space is given by $\Omega = \{(H, H), (H, T), (T, H), (T, T)\}$. We define the event that coin i comes up heads as $S_i = \{(v_1, v_2) \in \Omega \mid v_i = H\}$, with the corresponding probability function given by $\mathbb{P}[S_i] = \frac{1}{2}$. By independence, the probability that both coins were heads is $\mathbb{P}[S_1 \cap S_2] = \mathbb{P}[S_1]\mathbb{P}[S_2] = \frac{1}{4}$.

4. The Erdős-Rényi Random Graph

In this section, we discuss Paul Erdős and Alfréd Rényi's finding, that there is a single model of the countable random graph which is unique up to isomorphism. In general, a countable random graph is constructed by fixing a countably infinite vertex set V, from which edges are selected independently with probability $\frac{1}{2}$. We will define our probability space using this notion.

Our sample space Ω will be the set of all graphs on the vertex set $V = \mathbb{N} = \{1, 2, 3, \cdots\}$. Note that although the general notion of the random graph doesn't specify that the vertex set *be* $\mathbb{N}$, all of the calculations in this section can be replicated on arbitrary countably infinite vertex sets, so fixing V merely helps to simplify our sample space. Define

$$S_{v,w} = \{G(V, E) \in \Omega \mid \{v, w\} \in E\}$$

as the event that v and w are neighbors. Since we want to consider the probabilities of edge events, the set of events will be the smallest σ-algebra containing each $S_{v,w}$. For any two distinct vertices v and w, we define our probability measure by $\mathbb{P}[S_{v,w}] = \frac{1}{2}$, and specify that each $S_{v,w}$ is independent – that is, if we have $S_{v,w}$ and $S_{v',w'}$ as defined above such that $\{v, w\} \neq \{v', w'\}$, then

$$\mathbb{P}[S_{v,w} \cap S_{v',w'}] = \mathbb{P}[S_{v,w}]\mathbb{P}[S_{v',w'}].$$

Theorem 4.1 (Erdős-Rényi)**.** *There exists a unique graph Q such that a random graph G is isomorphic to Q with probability* 1*.*

Note that this theorem does not preclude the *existence* of non-Q countable graphs – the null and complete graphs are the most obvious examples. Additionally, it is not obvious that "uniqueness up to isomorphism" is even the type of property whose probability we can compute – especially without a model of Q already on hand. Thus, our proof will rely heavily on the extension property, which turns out to be a sufficient condition for demonstrating isomorphism. An outline of the proof below is given in Peter Cameron's "Random Graph Notes" [2], although the approach given here will be a bit more technical.

Definition 4.2. Let $G(V, E)$ be a graph. G has the *extension property* if for any finite disjoint subsets X and Y of the vertex set V, there exists some vertex $z \in V \setminus (X \cup Y)$ such that $z \sim x$ for all $x \in X$ and $z \not\sim y$ for all $y \in Y$. We will say for brevity that z *extends* X *and* Y.

Lemma 4.3. *A countable random graph has the extension property with probability* 1*.*

Proof. Let $G(V, E)$ be a countable random graph. Fix X and Y as finite, disjoint subsets of V, and consider an arbitrary vertex $z \in V \setminus (X \cup Y)$. According to the probability space defined above, denote

$$S_{X,Y,z} = \bigcap_{x \in X} S_{x,z} \cap \bigcap_{y \in Y} S_{y,z}^c$$

as the event that $z \sim x$ for all $x \in X$ and $z \not\sim y$ for all $y \in Y$. By event independence, we then see that:

$$\mathbb{P}[S_{X,Y,z}] = \prod_{x \in X} \mathbb{P}[S_{x,z}] \cdot \prod_{y \in Y} \mathbb{P}[s_{y,z}] = 2^{-|X|-|Y|}.$$

Let $S_{X,Y} = \bigcup_{z \in V \setminus (X \cup Y)} S_{X,Y,z}$ be the event that there exists some z extending X and Y. From set theory, we have

$$S^c_{X,Y} = \bigcap_{z \in V \setminus (X \cup Y)} S^c_{X,Y,z}.$$

Let q be the probability that a given vertex z does *not* extend X and Y, $q = \mathbb{P}[S^c_{X,Y,z}]$. From above, we see that $q = 1 - 2^{-|X|-|Y|}$, meaning that we have $0 < q < 1$ for any sizes of X and Y (so long as $X \cup Y \neq \emptyset$). Thus, if $z_1, \cdots, z_k$ are distinct vertices not contained in X and Y, the probability that none of them extend X and Y is q^k. Since $z_1, \cdots, z_k \in V \setminus (X \cup Y)$, we have

$$\bigcap_{z \in V \setminus (X,Y)} S^c_{X,Y,z} \subseteq S^c_{X,Y,z_1} \cap \cdots \cap S^c_{X,Y,z_k}.$$

By monotonicity, this gives us

$$\mathbb{P}[\bigcap_{z \in V \setminus (X,Y)} S^c_{X,Y,z}] \leq \mathbb{P}[S^c_{X,Y,z_1} \cap \cdots \cap S^c_{X,Y,z_k}]$$

for any choice of k and vertices $z_1, \cdots, z_k$. Note that as $k \to \infty, \mathbb{P}[S^c_{X,Y,z_1} \cap \cdots \cap S^c_{X,Y,z_k}] = q^k \to 0$, and so the probability $\mathbb{P}[S^c_{X,Y}]$ that X and Y are not extendable is always 0.

Let S be the event that G satisfies the extension property (i.e. every choice of X, Y is extendable), $S = \bigcap_{X,Y} S_{X,Y}$. Equivalently, we have $S^c = \bigcup_{X,Y} S^c_{X,Y}$. By the union bound, we see that

$$\mathbb{P}[S^c] \leq \sum_{X,Y} \mathbb{P}[S^c_{X,Y}].$$

From above, we know that $\mathbb{P}[S^c_{X,Y}] = 0$ for any choice of X and Y, so $\mathbb{P}[S^c] = \sum_{X,Y} \mathbb{P}[S^c_{X,Y}] = 0$ and $\mathbb{P}[S] = 1$. □

Lemma 4.4. *Any two countable graphs satisfying the extension property are isomorphic.*

Proof. Suppose $G(V, E)$ and $H(W, F)$ are two graphs satisfying the extension property. As one might deduce from the name, we can use this property to "extend" a bijective map to an additional vertex while preserving adjacencies. By assumption, the vertex sets V and W are countable, so we can enumerate them as $V = \{v_1, v_2, \cdots\}$ and $W = \{w_1, w_2, \cdots\}$. We can inductively construct a sequence of functions, each of which restricts to an isomorphism. Beginning with $f_0 = \emptyset$, suppose f_{n-1} has already been constructed. We wish to extend it to an additional vertex to construct f_n. Both vertex sets can be completely accounted for using a back-and-forth method.

If n is odd, let v_m be the lowest-indexed vertex in V which is not in the domain of f_{n-1}. Denoting this domain as $\mathrm{dom}(f_{n-1})$, let $X = \{v_i \in \mathrm{dom}(f_{n-1}) \mid v_i \sim v_m\}$ and $Y = \mathrm{dom}(f_{n-1}) \setminus X$. By assumption, $\mathrm{dom}(f_{n-1})$, and consequently X and Y are finite. We can then use the extension property on H to find the lowest-indexed $w_k \in W$ extending the images $f_{n-1}(X)$ and $f_{n-1}(Y)$, and assign $w_k = f_n(v_m)$.

If n is even, let w_m be the lowest-indexed vertex in W which is not in the image of f_{n-1}, which we will call $\mathrm{Im}(f_{n-1})$. Let $X = \{w_i \in \mathrm{Im}(f_{n-1}) \mid w_i \sim w_m\}$ and $Y = \mathrm{Im}(f_{n-1}) \setminus X$. By the extension property on G, we can find the lowest-indexed $v_k \in V$ extending the preimages $f^{-1}_{n-1}(X)$ and $f^{-1}_{n-1}(Y)$, and assign $w_m = f_n(v_k)$.

Now, let f be the union of the sequence of partial maps; by construction $\text{dom}(f) = V$ and $\text{Im}(f) = W$. The above steps tell us that f preserves all adjacencies, so f is the required isomorphism. □

Theorem 4.1 immediately follows from Lemma 4.3 and Lemma 4.4, thus concluding the proof. Since we have shown the random graph to be unique, from now on we will refer to the random graph as Q, and may assume the extension property on Q by default.

4.1. **A Deterministic Construction of Q.** Although notion of the random graph describes a probabilistically-generated structure, the uniqueness of Q allows us to model it deterministically as well. That is, any graph satisfying the extension property, regardless of construction, is a copy of the random graph Q. One such construction is that of the graph D on the vertex set $\mathbb{N}_0 = \{0, 1, 2, \cdots\}$, where adjacencies are determined by a number's binary expansion.

Definition 4.5. Let $n \in \mathbb{N}$. The *binary expansion* of n is given by

$$n = 2^k + a_{k-1}2^{k-1} + \cdots + a_0 2^0,$$

where $k \in \mathbb{N}_0$ and $a_i \in \{0, 1\}$. Additionally, the binary expansion of 0 is 0.

An implication of this definition is that every number in $\mathbb{N}_0$ can be uniquely expressed by its binary expansion, and no two numbers share a binary expansion.

Given the vertex set $\mathbb{N}_0$, we construct D by adding adjacencies as follows. Starting with the null graph on $\{0\}$, suppose we have established all adjacencies (and non-adjacencies) within the set $\{0, \cdots, n-1\}$. We then add the adjacencies between this set and n according to its binary expansion, $n = 2^k + a_{k-1}2^{k-1} + \cdots + a_0 2^0$. Taking $a_k = 1$, we say that $n \sim m$ if $a_m = 1$, and $n \not\sim m$ if $a_m = 0$.

To check for the extension property, fix finite, disjoint sets $X, Y \subset \mathbb{N}_0$; we want to show that there is some z extending them. Let $k = \max(X \cup Y) + 1$, and enumerate X in ascending order as $\{x_1, \cdots, x_m\}$. Then,

$$z = 2^k + (2^{x_m} + \cdots + 2^{x_1})$$

extends X and Y. Note that this k ensures that we're picking a z greater than all elements of X and Y, $z > 2^n$ for all $n \in X \cup Y$. In addition to ensuring that $z \in \mathbb{N}_0 \setminus (X \cup Y)$, this also guarantees that z is not a neighbor of any $y \in Y$, as all adjacencies within $\{0, \cdots, z\} \supset Y$ are already determined.

5. Some Corollaries to Erdős-Rényi

As shown in Section 4, the defining feature of the random graph is that it satisfies the extension property. The following proposition is a seemingly-stronger version of the extension property, which in fact follows directly from the property itself.

Proposition 5.1. *Let X and Y be be finite disjoint sets of vertices on the random graph $Q(V, E)$. Then the set*

$$Z = \{z \in V \setminus (X \cup Y) \mid z \sim x \text{ for all } x \in X; z \not\sim y \text{ for all } y \in Y\}$$

is infinite, and the induced subgraph on this set (i.e. the graph taking only the vertices in Z and their corresponding edges) is isomorphic to Q.

Proof. Fix finite, disjoint vertex subsets X and Y, and suppose for the sake of contradiction that Z is finite. By the extension property, we can find some vertex z' extending $X \cup Z$ and Y – that is, some $z' \in V \setminus (X \cup Y \cup Z)$ such that $z' \sim x$ for all $x \in X \cup Z$ and $z' \not\sim y$ for all $y \in Y$. We see that z' extends X and Y, so $z' \in Z$, but by the statement of the extension property we've already specified that $z' \notin Z$, so Z must be infinite.

To show that the induced subgraph on Z is isomorphic to Q, it is sufficient to demonstrate the extension property. Let X', Y' be finite disjoint subsets of Z. By the extension property on Q, there exists some $z' \in V \setminus ((X \cup X') \cup (Y \cup Y')$ extending $X \cup X'$ and $Y \cup Y'$. We see that $z' \in Z$ and z' extends X' and Y', thus proving our claim. □

A consequence of the above proposition is that the random graph resists certain types of finite changes.

Theorem 5.2 (Indestructibility)**.** *The result of any of the following operations on $Q(V, E)$ is isomorphic to Q:*

(a) Deleting a finite number of vertices;
(b) Adding or deleting a finite number of edges;
(c) Switching with respect to a finite number of vertices, i.e. for a finite vertex subset W, all edges involving some $v \in W$ are changed to non-edges and vice versa.

Proof. In all cases, it is sufficient to show that the extension property still holds. Denote the graph resulting from these operations as $Q'(V', E')$.

(a) Let W be the set of deleted vertices, such that the vertex set of Q' is $V' = V \setminus W$. If X and Y are finite, disjoint subsets of V', Proposition 5.1 tells us that since the set Z of vertices extending X and Y is infinite, we can find some $z \in Z \setminus A \subseteq W'$ extending X and Y.
(b) Let F be the set of ordered pairs corresponding to edges added to or deleted from E, $F = (E \setminus E') \cup (E' \setminus E)$. Since F is finite, it must be the case that the vertex set $W = \{v \mid \{v, w\} \in F\}$ is also finite, allowing us to to avoid any modified edges using the same method as in (a).
(c) Let W be the set of switched vertices, and consider finite disjoint vertex sets $X, Y \subset V'$. Applying the extension property on Q to $(X \setminus W) \cup (Y \cap W)$ and $(Y \setminus W) \cup (X \cap W)$ immediately gives us the extension property on Q'.

□

Note that there are certain finite changes that *do* destroy the isomorphism. For example, isolating a finite number of vertices causes the extension property to fail.

Another trait of Q is *partition regularity*, the notion that any finite partition of a set (or in this case, a graph) results in the original set. In the case of the random graph, this property is equivalent to the statement that any partition of a graph satisfying the extension property generates at least one part also satisfying the extension property.

Theorem 5.3 (Partition Regularity)**.** *If the vertex set of Q is partitioned into finitely many parts, then the the induced subgraph on one of these parts is isomorphic to Q.*

Proof. Suppose for the sake of contradiction that we can partition the vertex set as $V = V_1 \cup \cdots \cup V_n$, all parts pairwise disjoint, where the extension property fails on the induced subgraph on each $V_i, i = 1, \cdots, n$. This would mean that for each V_i, there exist finite disjoint sets $X_i, Y_i \subseteq V_i$ which are not extendable. Consider the sets $X = X_1 \cup \cdots \cup X_n$ and $Y = Y_1 \cup \cdots \cup Y_n$, finite and disjoint by assumption. If we take any $z \in V$, z must be in one of the parts V_i. Then z does not extend X_i and Y_i, and thus cannot extend X and Y, so the extension property fails on Q. □

Additionally, partition regularity is (almost) unique to Q:

Theorem 5.4. *The only partition-regular countable graphs (i.e. graphs with the property from Theorem 5.3) are the null graph, the complete graph, and the random graph Q.*

Proof. Suppose $G(V, E)$ is partition-regular, and is not null or complete. First, we can note that G cannot have any isolated vertices – for if they existed, we could partition them out and the part with no isolated vertices would have to be isomorphic to G. The same argument would also show that G has no vertices connected to all others.

Now, suppose that G is not isomorphic to Q. Then we can find sets X and Y for which extension fails; pick X and Y such that their union is minimal. Let $n = |X \cup Y|$. By above, G has no isolated vertices or vertices connected to all others, so $n > 1$ (as a single vertex could always be extended). This lets us partition $X \cup Y$ into two nonempty sets A and B. Now, let V_1 consist of A and all vertices in $V \setminus B$ which don't extend $X \cap A$ and $Y \cap A$, and let $V_2 = V \setminus V_1$. By assumption, all vertices in V_2 are either in B or extend $X \cap A$ and $Y \cap A$. Since extension fails for X and Y on G, all $v \in V_2$ fail to extend $X \cap B$ and $Y \cap B$. Because both A and B are nonempty, both $(X \cup Y) \cap A$ and $(X \cup Y) \cap B$ are smaller than $X \cup Y$, so the induced subgraphs on both V_1 and V_2 fail the extension property for sets smaller than X and Y. X and Y were assumed to be the smallest sets failing the extension property on G, so neither subgraph is isomorphic to G – thus, a contradiction. □

As we observed in the previous section, if two graphs have the extension property, we can construct an isomorphism between the two through a sequence of partial maps. It turns out that if only one of the graphs satisfies the extension property, we can still embed the graph that doesn't satisfy extension onto the one that does. That is, any finite or countable graph can be observed as a subgraph of Q.

Theorem 5.5 (Universality)**.** *Every finite or countable graph can be embedded as an induced subgraph of Q.*

Proof. Denote the random graph as $Q(V, E)$, and let $G(W, E)$ be a finite or countable graph. We can enumerate the vertex sets of Q and G as $V = \{v_1, v_2, \cdots\}$ and $W = \{w_1, w_2, \cdots\}$ respectively. Similar to the method in *Lemma 4.4*, we can construct a sequence of partial embeddings, starting from $f_0 = \emptyset$. Having constructed f_{n-1}, take the lowest-indexed $w_m \in W \setminus \text{dom}(f_{n-1})$ and assign as $f_n(w_k)$ the lowest-indexed $v_k \in V \setminus \text{Im}(f_{n-1})$ such that $v_k \sim f_{n-1}(w_i)$ if and only if $w_m \sim w_i$, for $w_i \in \text{dom}(f_{n-1})$. The union of these partial maps is our desired embedding. □

6. The 0-1 Law For Random Graphs

It happens that the extension property also gives us information about finite random graphs through the statement of the 0-1 law. In general, a *0-1 law* is a

statement that sentences in some formal language are either almost always true or almost always false on some class of finite structures. The 0-1 law for random graphs is concerned with whether sentences from first-order logic are true on finite random graphs. Specifically, the probability that a sentence is true is a function of the number of vertices in a graph – as the vertex set grows, this probability approaches either 0 or 1. Refining this statement gives us the following theorem:

Theorem 6.1 (0-1 Law for Random Graphs)**.** *Let ψ be a sentence from first-order logic. Denote $p_n(\psi)$ as the probability that ψ is true on a random graph with n vertices. If ψ is true on the infinite random graph, then $\lim_{n\to\infty} p_n(\psi) = 1$. Conversely, if $\neg\psi$ is true on the infinite random graph, then $\lim_{n\to\infty} p_n(\psi) = 0$.*

In order to make sense of this theorem, we need to learn some notions of first-order logic. A more comprehensive discussion of these concepts is given in chapter 1 of Marker's *Model Theory: An Introduction* [3].

To begin, we need an underlying structure, which we call a language.

Definition 6.2. A *language* $\mathcal{L}(\mathcal{F}, \mathcal{R}, \mathcal{C})$ is given by a set of function symbols $\mathcal{F}$, a set of relation symbols $\mathcal{R}$, and a set of constant symbols $\mathcal{C}$. Each function symbol $f \in \mathcal{F}$ and relation symbol $R \in \mathcal{R}$ specifies a number of arguments n_f or n_R respectively.

For our discussion today, we will be working with the language of graphs, which contains only a single binary relation $\sim$. Thus, we will focus on relational languages, i.e. those with no function symbols.

Definition 6.3. An *$\mathcal{L}$-structure* or *model* $\mathcal{M}$ is given by a nonempty set M which we call the *universe* of $\mathcal{M}$, functions $f^{\mathcal{M}} : M^{n_f} \to M$ corresponding to each $f \in \mathcal{F}$, sets $R^{\mathcal{M}} \subseteq M^{n_R}$ corresponding to each relation symbol $R \in \mathcal{R}$, and elements $c^{\mathcal{M}} \in M$ corresponding to each constant symbol $c \in \mathcal{C}$. We call $f^{\mathcal{M}}, R^{\mathcal{M}}$, and $c^{\mathcal{M}}$ the *interpretations* of the symbols in $\mathcal{L}$.

Essentially, an $\mathcal{L}$-structure is a specific example of the "class" of objects categorized by the language $\mathcal{L}$. For the language of graphs $\mathcal{G}$, the set of $\mathcal{G}$-structures is the simply set of structures equipped with a binary relation. This includes the set of all graphs (where this relation is irreflexive and symmetric), but also includes other structures with such a relation.

Definition 6.4. An *$\mathcal{L}$-term* (or simply, *term*) t is either a constant $c \in \mathcal{C}$, a variable x_i for $i = 1, 2, \cdots$, or a function of constants and variables, $t = f(t_1, \cdots, t_{n_f})$, where $f \in \mathcal{F}$ and each t_i is a term. Note that terms are finitely complex – that is, a term can only have finitely many subterms.

If $\mathcal{M}$ is an $\mathcal{L}$-structure and t is a term involving variables $x_{i_1}, \cdots, x_{i_m}$, we can interpret $t^{\mathcal{M}}$ as an m-ary function. Denote the list of arguments as $\overline{a} = (a_{i_1}, \cdots, a_{i_m}) \in M^m$. In relational languages, we can simply take $t^{\mathcal{M}}(\overline{a}) = c^{\mathcal{M}}$ or $t^{\mathcal{M}}(\overline{a}) = v_i$ based on the construction of t. For a more detailed explanation of how to interpret $\mathcal{L}$-terms, see [3].

Definition 6.5. An *atomic $\mathcal{L}$-formula* ϕ_0 associates constants and variable symbols either by equality, $\phi_0 = (t_1 = t_2)$, or by relation symbol, $\phi_0 = R(t_1, \cdots, t_{n_r})$ where each t_i is a term and $R \in \mathcal{R}$.

A *$\mathcal{L}$-formula* ϕ is a boolean combination of preexisting formulas, created by combining formulas using the "or" symbol $\vee$ or the "and" symbol $\wedge$, or by appending

the negation $\neg$ or quantifiers $\exists x$ or $\forall x$. (We additionally use $\psi \to \theta$ to abbreviate $\neg\psi \vee \theta$, and $\psi \leftrightarrow \theta$ for $(\psi \to \theta) \wedge (\theta \to \psi)$.) Thus, a formula ϕ is one of the following:

- ϕ is an atomic formula;
- $\phi = \neg\psi$ for a formula ψ;
- $\phi = \psi \wedge \theta$ or $\phi = \psi \vee \theta$ for formulas ψ and θ;
- $\phi = \exists x_i \psi$ or $\phi = \forall x_i \psi$ for a variable x_i and a formula ψ.

Also, note again that this definition implies finite complexity, i.e. a formula has finitely many subformulas.

Example 6.6. The set of graphs is described by the sentences

$$\forall x(\neg x \sim x) \quad \text{and} \quad \forall x \forall y (x \sim y) \to (y \sim x),$$

i.e. the relation $\sim$ is irreflexive and symmetric.

Remark 6.7. We say a variable x_i is *free* in a formula ϕ if it is not inside a quantifier $\exists x_i$ or $\forall x_i$; otherwise x_i is *bound*. A formula ϕ is a *sentence* if it has no free variables.

Definition 6.8. For a formula ϕ with m free variables, an $\mathcal{L}$-structure $\mathcal{M}$, and a list of arguments $\overline{a} = (a_{i_1}, \cdots, a_{i_m}) \in M^m$, we use $\mathcal{M} \models \phi(\overline{a})$ to denote when $\phi(\overline{a})$ is *true* on $\mathcal{M}$. (Equivalently, we can say $\mathcal{M}$ *satisfies* $\phi(\overline{a})$.) This is inductively defined on the structure of ϕ:

- If $\phi = (t_1 = t_2)$, then $\mathcal{M} \models \phi(\overline{a})$ if $t_1^{\mathcal{M}}(\overline{a}) = t_2^{\mathcal{M}}(\overline{a})$;
- If $\phi = R(t_1, \cdots, t_{n_R})$, then $\mathcal{M} \models \phi(\overline{a})$ if $(t_1^{\mathcal{M}}(\overline{a}), \cdots, t_{n_R}^{\mathcal{M}}(\overline{a})) \in R^{\mathcal{M}}$;
- If $\phi = \neg\psi$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \not\models \psi(\overline{a})$;
- If $\phi = \psi \wedge \theta$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a})$ and $\mathcal{M} \models \theta(\overline{a})$;
- If $\phi = \psi \vee \theta$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a})$ or $\mathcal{M} \models \theta(\overline{a})$;
- If $\phi = \exists x_i \psi$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a}, b)$ for some $b \in M$;
- If $\phi = \forall x_i \psi$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a}, b)$ for all $b \in M$.

If ϕ has no free variables, we may simply write $\mathcal{M} \models \phi$.

Definition 6.9. A theory T is a set of $\mathcal{L}$-formulas. We say that a structure $\mathcal{M}$ *is a model of* T and write $\mathcal{M} \models T$ if $\mathcal{M} \models \phi$ for all $\phi \in T$. T is *satisfiable* if it has a model; otherwise, T is *inconsistent*.

Having defined what it means for a sentence ψ to be true on a model $\mathcal{M}$, we may now define our probability measure for finite graphs. We obtain a finite random graph G by fixing a vertex set and selecting edges independently with probability $\frac{1}{2}$. Fix $n \in \mathbb{N}$. Considering the sample space as the set of all graphs on $\{1, 2 \cdots, n\}$, we assign each graph an equal likelihood to "appear". Thus, the probability of picking a graph where a formula ψ is true is

$$p_n(\psi) := \frac{|\{G(V,E) \mid G \models \psi\}|}{2^{\binom{n}{2}}}.$$

We will discuss two methods of proving the 0-1 law – one using the compactness theorem for first-order logic, and the other utilizing the Ehrenfeucht–Fraïssé game, a back-and-forth technique for determining elementary equivalence. For either method, it is sufficient to consider the extension property as a theory consisting of a sequence of analogous and increasingly-stronger sentences, $\Phi = \{\phi_1, \phi_2, \cdots\}$.

Define

$$\phi_n = \forall x_1, \cdots, x_n, y_1, \cdots, y_n \exists z \bigwedge_{i,j=1}^{n} \neg(x_i = y_i) \to \bigwedge_{i=1}^{n} z \sim x_i \wedge \bigwedge_{i=1}^{n} \neg(z \sim y_i).$$

Essentially, ϕ_n asserts that the extension property holds for sets X and Y of size n.

Remark 6.10. At first glance, Φ seems weaker than the statement of the extension property from Section 4, which does not specify that X and Y be the same size. If X and Y are different sizes, we can simply add elements to the smaller set – for if the larger sets are extendable, X and Y clearly are as well. That is to say, if a graph satisfies ϕ_n, all sets of size *up to* n are extendable; so a graph modeling Φ necessarily satisfies the extension property. We can thus treat Φ and the extension property as equivalent conditions.

7. The Compactness Theorem

In first-order logic, *compactness* is the notion that if every finite subset of a theory is satisfiable, the theory itself must also be satisfiable. This section's proof of the 0-1 law utilizes the compactness theorem as well as the Löwenheim-Skolem theorem, which allows us to restrict our purview to countable models. For this section, we reference chapter 2 of Marker's *Model Theory: An Introduction* [3].

Theorem 7.1 (Compactness)**.** *A theory T is satisfiable if and only if T is finitely satisfiable, i.e. every finite subset of T is satisfiable.*

Theorem 7.2 (Löwenheim-Skolem)**.** *Let T be a countable theory. If T is satisfiable, then T has a countable model.*

Both of these proofs are rather involved; they can be found in chapter 2 of [3]. In order to tie compactness to the theory Φ from the previous section, we need the following proposition.

Proposition 7.3. *For a theory T and a sentence ϕ, ϕ is true in every model of T if and only if $T \cup \{\neg\phi\}$ is inconsistent.*

Proof. Looking to Definition 6.8, the proof in either direction follows from the implication that for any model $\mathcal{M}$, exactly one of $\mathcal{M} \models \phi$ and $\mathcal{M} \models \neg\phi$ can be true. If $T \vdash \phi$ then any model of T (and by extension of ϕ) cannot satisfy $\neg\phi$; conversely if $T \cup \{\neg\phi\}$ has no model then any model of T must satisfy ϕ. □

Even when a formula ϕ is not an element of a theory T, we'd still like to have a notion of when T "naturally leads to" ϕ. We will write $T \vdash \phi$ to denote when a formula ϕ is true in every model of a theory T.

Corollary 7.4. *Let T be a countable theory, and suppose we have a formula ϕ such that for every countable model $\mathcal{M}$ of T, we have $\mathcal{M} \models \phi$. Then there exists a finite subset $F \subseteq T$ such that $F \vdash \phi$.*

Proof. Suppose for the sake of contradiction that for all finite $F \subseteq T, F \nvdash \phi$. By Proposition 7.3, this is equivalent to saying that each $F \cup \{\neg\phi\}$ is satisfiable. Thus, $T \cup \{\neg\phi\}$ is satisfiable, and so by Löwenheim-Skolem, $T \cup \{\neg\phi\}$ has a countable model. If $\mathcal{M}$ is such a model, then we have $\mathcal{M} \models T$ and $\mathcal{M} \not\models \phi$, contradicting the initial assumption. □

With these, we may finally prove the 0-1 law.

Proof of Theorem 6.1 using compactness. The proof of this fact is very similar in form to Lemma 4.3. As shown in Section 4, Q is the only countable graph satisfying the extension property, and is thus the only graph satisfying Φ. Then, for a sentence ψ with $Q \models \psi$, we can apply Corollary 7.4 to find a finite subset $F \subset \Phi$ with $F \vdash \psi$. Observe that as shown in *Remark* 6.10, each $\phi_k \in \Phi$ is a stronger version of the sentences earlier in the sequence, $\{\phi_k\} \vdash \phi_m$ for $m \leq k$. Then, if ϕ_k is the highest-indexed element of F, we have $\{\phi_k\} \vdash F$, and so $\{\phi_k\} \vdash \psi$. By how we defined p_n, we see that $p_n(\phi_k) \leq p_n(\psi)$. Thus, it is sufficient to show that $\lim_{n\to\infty} p_n(\phi_k) = 1$.

Assuming G is a graph with $n \geq 2k+1$ vertices, fix $x_1, \cdots, x_k, y_1, \cdots, y_k$, and let z be a vertex distinct from all of these. Since edges are selected independently with probability $\frac{1}{2}$, we have

$$p_n\left(\bigwedge_{i=1}^{k} z \sim x_i \wedge \bigwedge_{i=1}^{k} \neg z \sim y_i\right) = 2^{-2k}.$$

Accordingly, the probability that z fails to extend $x_1, \cdots, x_k, y_1, \cdots, y_k$ is

$$q = p_n\left(\neg\left(\bigwedge_{i=1}^{k} z \sim x_i \wedge \bigwedge_{i=1}^{k} \neg z \sim y_i\right)\right) = 1 - 2^{-2k}.$$

By edge independence, the probability that

$$G \models \neg\exists z\left(\bigwedge_{i=1}^{k} z \sim x_i \wedge \bigwedge_{i=1}^{k} \neg z \sim y_i\right)$$

is q^{n-2k}. Observe that $\neg\phi_k$ is the union of $\neg\exists z\left(\bigwedge_{i=1}^{k} z \sim x_i \wedge \bigwedge_{i=1}^{k} \neg z \sim y_i\right)$ over all choices of $x_1, \cdots, x_k, y_1, \cdots, y_k$. Let N be the number of pairs of disjoint subsets of size k, i.e. the number of ways to pick $x_1, \cdots, x_k, y_1, \cdots, y_k$. Using the upper bound $N \leq \binom{n}{k}^2 < n^{2k}$, we can apply the union bound to see that

$$p_n(\neg\phi_k) \leq Nq^{n-2k} < n^{2k}q^{n-2k}.$$

Since $0 \leq q < 1$, computing the limit gives us

$$\lim_{n\to\infty} p_n(\neg\phi_k) = \lim_{n\to\infty} n^{2k}q^{n-2k} = 0.$$

By the first part of the proof, this gives us our desired result of $\lim_{n\to\infty} p_n(\psi) = 1$. Conversely, if $Q \models \neg\psi$, then $\lim_{n\to\infty} p_n(\neg\psi) = 1$, so $\lim_{n\to\infty} p_n(\psi) = 0$. □

8. Ehrenfeucht–Fraïssé Games

As we saw in the previous section, the proof of the 0-1 law follows rather nicely from compactness. However, there is no indication of how quickly the probability of a given sentence ψ stabilizes – although the proof utilizing compactness does use some ϕ_n as a lower bound on the likelihood of ψ, we are given no indication as to what this critical ϕ_n could be. The Ehrenfeucht–Fraïssé game fixes this by associating ψ with some specific ϕ_n corresponding to its quantifier rank (which we wll define promptly). For further reference, see Jouko Väänänen's *Models and Games* [4].

Definition 8.1. Let ϕ be an $\mathcal{L}$-formula. We define its *quantifier rank* $QR(\phi)$ in the following way:

- If ϕ is atomic, then $QR(\phi) = 0$;
- If $\phi = \neg\psi$, then $QR(\phi) = QR(\psi)$;
- If $\phi = \psi \wedge \theta$ or $\phi = \psi \vee \theta$, then $QR(\phi) = \max\{QR(\psi), QR(\theta)\}$;
- If $\phi = \exists v_i \psi$ or $\phi = \forall v_i \psi$, then $QR(\phi) = QR(\psi) + 1$.

Note that by this definition, the quantifier rank of a formula is only increased by (as the name suggests) attaching quantifiers.

Example 8.2. The sentence indicating that any two vertices are connected by a path of length ≤ 2 is given by

$$\forall x \forall y (x \sim y) \vee \exists z (z \sim x \wedge z \sim y)$$

and has quantifier rank 3.

Definition 8.3. We say that two structures $\mathcal{M}$ and $\mathcal{N}$ are *elementarily equivalent up to rank* n if they satisfy the same formulas of quantifier rank $\leq n$, and denote this as $\mathcal{M} \equiv_n \mathcal{N}$. If $\mathcal{M}$ and $\mathcal{N}$ satisfy the same formulas of any rank, we simply call them *elementarily equivalent* and write $\mathcal{M} \equiv \mathcal{N}$.

The Ehrenfeucht–Fraïssé game is played by two players (whom we will call Player I and Player II) on two graphs for a predetermined number of rounds. For the sake of simplicity, we consider the graphs to be disjoint. We denote the Ehrenfeucht–Fraïssé game played on graphs $G(V, E)$ and $H(W, F)$ with n rounds as $EF_n(G, H)$. Player I's goal is to demonstrate a difference between the two graphs, while Player II wants to show that they are the same (up to isomorphism).

Supposing that we've just finished round $k-1$, we enumerate the vertices played in G as $v_1, \cdots, v_{k-1}$ and those played in H as $w_1, \cdots, w_{k-1}$. Round k plays out in the following way: At the beginning of the round, Player I plays a vertex on one of the two graphs – we label this vertex v_k if it's played on G and w_k if on H. To this, Player II must respond with a vertex on the other graph, again labeled v_k if played on G and w_k if on H. We call vertices v_k and w_k which are played in the same round *corresponding vertices*. Also, we'll denote the vertices played in the rounds up to round k as the ordered k-ples $V_k = (v_1, \cdots, v_k)$ and $W_k = (w_1, \cdots, w_k)$, and refer to the state of play as *position* (V_k, W_k).

At the end of round n, we observe the played vertices $V_n = (v_1, \cdots, v_n)$ and $W_n = (w_1, \cdots, w_n)$ and check for two conditions:

(1) $v_i = v_j$ if and only if $w_i = w_j$;
(2) $v_i \sim v_j$ if and only if $w_i \sim w_j$.

If both conditions are satisfied, Player II wins; otherwise, Player I wins. In essence, this process seeks to construct n partial isomorphisms between G and H, stopping either when the current map cannot be extended or upon successfully constructing n maps.

We can characterize a given Ehrenfeucht–Fraïssé game by the strategies each player follows. A *strategy* is a function τ prescribing a player's next move based on what vertices have been played so far. A strategy τ is a *winning strategy* if, for any sequence of vertices chosen by the opponent, a player can win by following τ.

Example 8.4. Suppose $G(V, E)$ and $H(W, F)$ are both graphs with 4 vertices. G is a 4-cycle, while H is a path of length 4. Player II has a winning strategy on

$EF_2(G, H)$. Suppose round 1 has concluded, and Player I plays $v_2 \in V$. Then, Player II picks $w_2 \in W$ such that $w_1 \sim w_2$ if and only if $v_1 \sim v_2$. (The same goes for if, instead, Player I plays $w_2 \in W$.)

On the other side, Player I has a winning strategy on $EF_3(G, H)$. In the first two rounds, Player I can play w_1 and w_2 on opposite ends of graph H. In order to avoid losing immediately, Player II would have to play opposite corners of graph G. Then, Player I can play either of the unplayed vertices of H – suppose he plays $w_3 \in W$ such that $w_3 \sim w_1$ and $w_3 \not\sim w_2$. Any choice of $v_3 \in V \setminus \{v_1, v_2\}$ results in $v_3 \sim v_1$ and $v_3 \sim v_2$.

The following proposition will be helpful for proving our main result. A proof is given in chapter 4 of [4].

Proposition 8.5. *For a given $n \in \mathbb{N}$ and set of variables $x_1, \cdots, x_n$, there are finitely many formulas of the language of graphs of quantifier rank $< n$ with free variables $x_1, \cdots, x_n$ up to logical equivalence.*

Theorem 8.6 (Ehrenfeucht–Fraïssé)**.** *Let G and H be graphs. The following are equivalent:*

(1) G and H satisfy the same formulas up to quantifier rank n;
(2) Player II has a winning strategy on $EF_n(G, H)$.

Proof. $(1 \Rightarrow 2)$ Suppose G and H satisfy the same formulas up to quantifier rank n. We claim that there exists a strategy for Player II such that if $k \leq n$, then after round k, G and H satisfy the same formulas of quantifier rank $\leq n-k$. In particular, if $QR(\phi) \leq n - k$, then $G \models \phi(\overline{v})$ whenever $H \models \phi(\overline{w})$, where $\overline{v}$ and $\overline{w}$ are lists of arguments containing corresponding vertices. At the beginning of the game, $V_0 = W_0 = \emptyset$, so our claim is true by assumption. Suppose we are about to begin round k, with the game's current position being (V_{k-1}, W_{k-1}). Without loss of generality, suppose Player I plays vertex $a = v_k$ on G. (The case where Player I plays on H is symmetrical.) Let $\overline{v}_i$ denote a list of arguments consisting of previously played variables on G. By Proposition 8.5, we may enumerate the formulas of quantifier rank $< n - k$ describing previously played variables as $\psi_1(\overline{v}_1, x), \cdots, \psi_m(\overline{v}_m, x)$. We then observe that

$$\phi = \bigwedge_{\substack{0 \leq i \leq m \\ G \models \psi_i(a)}} \psi_i(\overline{v}_i, x) \wedge \bigwedge_{\substack{0 \leq i \leq m \\ G \not\models \psi_i(a)}} \neg\psi_i(\overline{v}_i, x)$$

is a formula of quantifier rank $< n - k$, so $\exists x\phi$ has quantifier rank $\leq n - k$. Since a acts as a witness to ϕ on G, we have $G \models \exists x\phi$; elementary equivalence then gives us $H \models \exists x\phi$ as well. Let $b \in W$ be a witness to $\exists x\phi$ on H; we can thus let Player II play $w_k = b$, proving the inductive hypothesis.

$(2 \Rightarrow 1)$ Suppose Player II has a winning strategy on $EF_n(G, H)$. We make a stronger claim for $k \leq n$: If Player II has a winning strategy with (V_{n-k}, W_{n-k}) having already been placed, then for all formulas ϕ of quantifier rank $\leq k$, $G \models \phi$ if and only if $H \models \phi$. Again, note that in particular, $G \models \phi(\overline{v})$ whenever $H \models \phi(\overline{w})$, for lists of arguments $\overline{v}$ and $\overline{w}$ containing corresponding vertices.

Suppose the claim is true up to $k - 1$. To check for k, suppose $EF_n(G, H)$ is in round $n - k$, and Player II has a winning strategy on (V_{n-k}, W_{n-k}). If ϕ is atomic, then the claim is demonstrated by the fact that Player II is currently in a winning position. Otherwise, suppose ϕ is a formula of quantifier rank k.

If $\phi = \exists x\psi$, then ψ is a formula of quantifier rank $k-1$. Suppose without loss of generality that $G \models \exists x\psi(x)$, and let Player I play $a = v_k \in V$ such that $G \models \psi(a)$. Since we've established that Player II has a winning strategy given (V_{n-k}, W_{n-k}), she can play $b = w_k \in W$ accordingly. Now, we are in round $n-k+1$, and Player II is still in a winning position; thus, we may apply the inductive hypothesis to ψ. Since $G \models \psi(a)$, we also have $H \models \psi(b)$, as a and b are corresponding vertices. Since b is a witness to $\exists x\psi$, we must have had $H \models \phi$ in round $n-k$.

If ϕ is any of $\neg\psi, \psi \wedge \theta, \psi \vee \theta$, or $\forall x\psi$, the claim follows from the case for $\exists x\psi$. In particular, the first three cases are directly given by whether the claim holds on the subformula(s), while $\forall x\psi$ is logically equivalent to $\neg\exists x\neg\psi$. □

Proposition 8.7. *Let G and H be graphs. If both G and H satisfy ϕ_n, then Player II has a winning strategy on $EF_n(G, H)$.*

Proof. Suppose without loss of generality that $EF_n(G, H)$ is in position (V_{k-1}, W_{k-1}) and Player I has just played $v_k \in V$. As we've observed, a graph satisfying ϕ_n in fact satisfies the extension property for sets X and Y up to size n. On graph H, we can let $X = \{w_i \mid v_i \sim v_k\}$ and $Y = \{w_i \mid v_i \not\sim v_k\}$; Player II can then select as w_k some vertex extending these. □

Proof of Theorem 6.1 using the Ehrenfeucht–Fraïssé game. Let ψ be a sentence with quantifier rank k with $Q \models \psi$. If G is a graph satisfying ϕ_k, we know by Proposition 8.7 that Player II wins on $EF_k(G, Q)$. By Theorem 8.6, this is equivalent to the assertion that $G \equiv_k Q$, so $G \models \psi$. Because ψ is true in every model of ϕ_k, we thus get $\{\phi_k\} \vdash \psi$. From the proof of the 0-1 law using compactness, we thus have

$$\lim_{n\to\infty} p_n(\phi_k) = \lim_{n\to\infty} p_n(\psi) = 1.$$

□

Acknowledgments

I'd like to thank my mentor Miles Kretschmer for his guidance throughout the REU. Model theory is incredibly new to me, and it's thanks to his patience and excellent teaching that I was able to learn it. I also want to thank Peter May for his hard work organizing the REU, through which I was introduced to so many new and interesting areas of math.

References

[1] Richard F. Bass. *Real Analysis For Graduate Students, Version 4.2*. CreateSpace, 2020.

[2] Peter J. Cameron. *The random graph*. https://arxiv.org/abs/1301.7544. 2013. arXiv: 1301.7544 [math.CO].

[3] David Marker. *Model Theory: An Introduction*. Springer, 2002.

[4] Jouko Väänänen. *Models and Games*. English. Cambridge Studies in Advanced Mathematics 132. United Kingdom: Cambridge University Press, May 2011. ISBN: 9780521518123.