Zero-One Laws, Random Graphs, and Fraïssé Limits

Joshua Horowitz

April 24, 2008

1 Introduction

While at first glance the fields of logic and probability may appear as immiscible as oil and water, a closer look reveals that probabilistic questions about logic often have surprisingly elegant and powerful answers. For instance, one might wonder "What is the probability that a certain sentence in a certain logic holds for a randomly chosen structure?" Though in many cases this question will turn out to be hopelessly complex, it turns out that some logics in some situations are not powerful enough to define properties with any probabilities other than 0 or 1. Results of this form, known as zero-one laws, are of central importance in the study of this fascinating collision between logic and randomness.

In this paper, we will begin by proving the zero-one law for first-order logic over graphs, using an ingenious construction known as the random graph. We will then use a related but much more general construction known as a Fraïssé limit in order to prove the zero-one law for first-order logic over the structures of any purely relational vocabulary.

2 Zero-One Laws and Random Graphs

2.1 Random Structures

Both for simplicity's sake and for more important reasons which will become apparent, the structures we are considering will only have relations, not functions or constants. That is, our vocabulary σ will always be *purely relational*.

A property $\mathcal{P}$ of finite σ-structures is just a class of finite σ-structures closed under isomorphism (so that relabeling a structure does not change whether it is in the class). The sort of question we will be concerned with is "What is the probability that a randomly chosen finite structure is in $\mathcal{P}$?" A problem with this question immediately jumps out at us: there are an infinite number of finite structures and there could easily be an infinite number of finite structures in $\mathcal{P}$. This means that we cannot express the probability we are looking for as a proportion in the obvious way:

$$\frac{\text{number of finite structures in } \mathcal{P}}{\text{number of finite structures}} = \frac{\infty}{\infty}.$$

It is clear that we need to be more careful about our definitions, in order to make sure that our questions make sense and that we know what we are talking about.

Choosing a random structure is not so hard if we restrict our consideration to σ-structures with universe some fixed finite set A, as there are only a finite number of structures on any finite set. We could thus define the probability that a randomly chosen structure on a fixed finite set A is in $\mathcal{P}$. But since $\mathcal{P}$ is closed under isomorphism, the only thing which distinguishes A is its size $n = \#(A)$. We might as well just use $A = \{0, \dots, n-1\}$ as our canonical universe. We denote the set of σ-structures on $\{0, \dots, n-1\}$ by $\textsc{Struct}_n[\sigma]$ and define the probability of $\mathcal{P}$ on σ-structures of size n as the ratio

$$\mu_n(\mathcal{P}) := \frac{\#\left(\textsc{Struct}_n[\sigma] \cap \mathcal{P}\right)}{\#\left(\textsc{Struct}_n[\sigma]\right)}.$$

As mentioned, we can't generalize this definition to the class $\textsc{Struct}[\sigma]$ of all finite σ-structures (since there are infinitely many of these), but we can approximate this concept by seeing what happens to $\mu_n(\mathcal{P})$ as n gets bigger and bigger. In particular, we define the asymptotic probability of $\mathcal{P}$ as

$$\mu(\mathcal{P}) := \lim_{n\to\infty} \mu_n(\mathcal{P}).$$

Note that we can also relativize these definitions to some specific class of structures $\mathcal{C}$, and find the asymptotic probability that a structure in the class $\mathcal{C}$ has property $\mathcal{P}$:

$$\mu_n(\mathcal{P} \mid \mathcal{C}) := \frac{\#\left(\textsc{Struct}_n[\sigma] \cap \mathcal{C} \cap \mathcal{P}\right)}{\#\left(\textsc{Struct}_n[\sigma] \cap \mathcal{C}\right)},$$
$$\mu(\mathcal{P} \mid \mathcal{C}) := \lim_{n\to\infty} \mu_n(\mathcal{P} \mid \mathcal{C}).$$

An especially common example of this is taking $\mathcal{C}$ to be the class $\mathcal{G}$ of graphs. For our purposes, a graph is an $\{E\}$-structure, with E interpreted as a symmetric, irreflexive binary relation. As defined, $\mu_n(\mathcal{P} \mid \mathcal{G})$ is the ratio of the sizes of two particular sets of structures, but we can think of it in more intuitive probabilistic terms. This is because, for any two vertices, exactly half of the (labeled) graphs on n vertices have the two edges joined by an edge and exactly half of them do not. In fact, for any selection of m distinct pairs of vertices, the set of graphs on n vertices can be split into 2^m equally-sized classes based on which of the m pairs they connect by an edge and which they do not. This means that we can compute $\mu_n(\mathcal{P} \mid \mathcal{G})$ by considering a random graph where there is a $\frac{1}{2}$ chance that there is an edge between any two given vertices, and these events are statistically independent. $\mu_n(\mathcal{P} \mid \mathcal{G})$ is just the probability that such a random graph has property $\mathcal{P}$.

Let's look at some possible behaviors these probabilities can exhibit:

- Take $\mathcal{P}_{\text{triangle}}$ to be the property of containing a triangle. If $n > 3k$ and $G = \{v_0, \ldots, v_{n-1}\}$, we can bound $\mu_n(\mathcal{P}_{\text{triangle}} \mid \mathcal{G})$ below by the probability that one of the sets of vertices $\{v_0, v_1, v_2\}, \ldots, \{v_{3k-3}, v_{3k-2}, v_{3k-1}\}$ forms a triangle. This probability is just $1 - \left(1 - \frac{1}{2^3}\right)^k = 1 - \left(\frac{7}{8}\right)^k$. As $n \to \infty$, we can take $k \to \infty$, and this lower bound tends toward 1 as $k \to \infty$, so $\mu(\mathcal{P}_{\text{triangle}} \mid \mathcal{G}) = 1$. We say that this property holds "almost surely".

- Take $\mathcal{P}_{\text{isolated vertex}}$ to be the property of having an isolated vertex. The property of any given vertex being isolated is $1/2^{n-1}$, so $\mu_n(\mathcal{P}_{\text{isolated vertex}} \mid \mathcal{G})$ is bounded above by $n/2^{n-1}$. Since this tends toward 0 as $n \to \infty$, we have $\mu(\mathcal{P}_{\text{isolated vertex}} \mid \mathcal{G}) = 0$. We say that this property holds "almost never".

- Take $\mathcal{P}_{\text{even edges}}$ to be the property of having an even number of edges. If $n(n-1)/2$ (the number of possible edges) is odd, $\mu_n(\mathcal{P}_{\text{even edges}} \mid \mathcal{G})$ is exactly $\frac{1}{2}$, and if $n(n-1)/2$ is even, then the probability is exactly $\frac{1}{2}$ once you disregard the possibility that the number of edges is exactly $n(n-1)/4$, which becomes less and less likely as $n \to \infty$. So $\mu(\mathcal{P}_{\text{even edges}} \mid \mathcal{G}) = \frac{1}{2}$; this property holds neither almost surely nor almost never.

- Take $\mathcal{P}_{\text{even vertices}}$ to be the property of having an even number of vertices. As $n \to \infty$, $\mu_n(\mathcal{P}_{\text{even vertices}} \mid \mathcal{G})$ jumps back and forth between 0 (n odd) and 1 (n even). Thus, the limit $\mu(\mathcal{P}_{\text{even vertices}} \mid \mathcal{G})$ does not even exist.

2.2 Zero-One Laws

Readers who know a thing or two about the expressive power of first-order logic will notice something interesting about the above four examples. The first two properties, $\mathcal{P}_{\text{triangle}}$ and $\mathcal{P}_{\text{isolated vertex}}$, are those which are definable in first-order logic. They are also the ones with the relatively uninteresting asymptotic probabilities of 1 and 0. In order to have asymptotic probabilities other than 1 or 0, it seems that the properties $\mathcal{P}_{\text{even edges}}$ and $\mathcal{P}_{\text{even vertices}}$ must demand a sort of counting which is not first-order definable.

It turns out that $\mu(\mathcal{P}_{\text{triangle}} \mid \mathcal{G}) = 1$ and $\mu(\mathcal{P}_{\text{isolated vertex}} \mid \mathcal{G}) = 0$ are just two particular cases of a general theorem called the zero-one law for first-order logic over graphs:

Theorem 2.1. *If the property $\mathcal{P}$ is first-order definable over graphs, then it must hold either almost surely or almost never, that is, $\mu(\mathcal{P} \mid \mathcal{G})$ must be either 1 or 0.*

This is just one example of a great number of similar zero-one laws which have been discovered. The general structure of a zero-one law is expressed in the following definition.

Definition 2.2. *A logic $\mathcal{L}$ has the* zero-one law *over a class $\mathcal{C}$ if for every property $\mathcal{P}$ definable in $\mathcal{L}$ over $\mathcal{C}$, $\mu(\mathcal{P} \mid \mathcal{C})$ is either 1 or 0.*

To prove that first-order logic has the zero-one law in some particular case, we will use the technique suggested by the following proposition.

Proposition 2.3. *Suppose there is an theory T of sentences such that 1. every sentence in T holds almost surely among the structures in $\mathcal{C}$ and 2. T is complete (that is, for every sentence ϕ, $T \models \phi$ or $T \models \neg\phi$). Then $\mathcal{L}$ has the zero-one law over $\mathcal{C}$.*

Proof. Suppose we are given some sentence ϕ. By completeness of T, either $T \models \phi$ or $T \models \neg\phi$. Suppose $T \models \phi$. Then ϕ follows from some finite number of sentences $\psi_1, \ldots, \psi_s \in T$ (by compactness). But each ψ_i holds almost surely among $\mathcal{C}$, so ϕ holds almost surely among $\mathcal{C}$. If, on the other hand, $T \models \neg\phi$, then by the above argument $\neg\phi$ holds almost surely, so ϕ holds almost never. □

We will use this technique to prove the zero-one law for first-order logic over graphs. The theory T used will consist of sentences known as *extension axioms*. Later, we will use the technique in conjunction with a construction known as the Fraïssé limit in order to prove the zero-one law for first-order logic over general structures.

Note that these results, and zero-one laws in general, rely on our signature being purely relational. For instance, in the theory of the random graph plus two constants a and b, we have the sentence $E(a, b)$ which has probability converging to $\frac{1}{2} \neq 0, 1$.

2.3 Extension Axioms

Extension axioms are certain sentences about graphs which are in one sense powerful and in another sense ubiquitous. They are powerful because, taken together, they imply that the graph contains every finite graph as a substructure (in fact, they show even more than this, as we shall see). On the other hand, they are ubiquitous, because as the number of vertices goes to infinity, the probability of any given extension axiom being true goes to 1. It is this twofold nature which allows us to use extension axioms to prove zero-one laws.

For $k, l \geq 0$, the extension axiom $EA_{k,l}$ says that given $k + l$ distinct vertices, a new vertex can be found which is adjacent to the first k and not adjacent to the last l. It can be expressed in first-order logic.

Definition 2.4. *The extension axiom $EA_{k,l}$ is the sentence*

$$\forall x_1 \cdots \forall x_{k+l} \left[\left(\bigwedge_{i \neq j} x_i \neq x_j \right) \Rightarrow \exists y \left(\bigwedge_i \left\{ \begin{array}{ll} E(x_i, y) & i \leq k \\ \neg E(x_i, y) \wedge x_i \neq y & i > k \end{array} \right\} \right) \right].$$

The theory EA is the set of all extension axioms: $EA = \{EA_{k,l} \mid k, l \geq 0\}$

Our theory T to use in Proposition 2.3 will be EA. To derive a zero-one law from EA, we have two things to prove. First, we must show that the elements of EA (the extension axioms) hold almost surely. After that, we must show that EA is complete. The first part is just some basic probability:

Proposition 2.5. $\mu(EA_{k,l} \mid \mathcal{G}) = 1$

Proof. We will prove that $\mu(\neg EA_{k,l} \mid \mathcal{G}) = 0$. The sentence $\neg EA_{k,l}$ says that there exist some $k + l$ distinct vertices such that no $(k + l + 1)^{\text{th}}$ vertex exists which is connected to the first k but not the last l. We want to prove that the probability of this goes to zero as the number of vertices in the graph goes to infinity.

Fix for a moment any $x_1, \ldots, x_{k+l}$. For each candidate y, there is a $1/2^{k+l}$ chance that it will have the correct connections to the x_i's. So the probability that none of the $n - k - l$ candidates will have the correct connections is exactly $(1 - 1/2^{k+l})^{n-k-l}$.

There are $\frac{n!}{(n-k-l)!}$ ways to pick the $x_1, \ldots, x_{k+l}$. Each one has a $(1 - 1/2^{k+l})^{n-k-l}$ probability of being the single failure we need to demonstrate $\neg EA_{k,l}$. The situation which would maximize the probability of there being at least one failure among all the possible choices of x's is one in which the failures never occur simultaneously (which would be inefficient overlap, if you're trying to maximize the probability of there being at least one). In this worst-case scenario, the probability of there being at least one failure is $\frac{n!}{(n-k-l)!}(1 - 1/2^{k+l})^{n-k-l} = \mathcal{O}\left(n^{k+l}(1 - 1/2^{k+l})^n\right)$. This upper bound tends towards 0 as $n \to \infty$, so $\mu(\neg EA_{k,l} \mid \mathcal{G}) = 0$, so $\mu(EA_{k,l} \mid \mathcal{G}) = 1$. □

Now we must prove completeness of EA. As one step in doing this, we will construct a model for EA called the *random graph*.

2.4 The Random Graph

Since EA lets us build up subgraphs of arbitrary size within its models, it certainly cannot have any finite models. The smallest size we can hope for is countable (that is, countably infinite). The random graph is such a countable model.

Hopefully not confusingly, we will construct the "random" graph by a deterministic process. Letting $[i]_j$ denote the j^{th} bit in the binary expansion of i, we define:

Definition 2.6. *The* random graph $\mathfrak{RG}$ *has vertices* $\{v_0, v_1, v_2, \ldots\}$*, with an edge between* v_i *and* v_j *iff* $[i]_j = 1$ *or* $[j]_i = 1$.

It is worth mentioning that there is, in fact, a good reason to call this the "random graph". If you build up a countably infinite graph by adding vertices one at a time, attaching each new vertex to any given old vertex with probability $\frac{1}{2}$, the resulting graph will be the random graph with probability 1.

As promised, $\mathfrak{RG}$ will form a model for EA.

Proposition 2.7. $\mathfrak{RG} \models EA$

Proof. We must verify that $\mathfrak{RG} \models EA_{k,l}$ for each pair $k.l$. So fix some k, l and suppose we are given disjoint $K, L \subseteq \{v_0, v_1, v_2, \ldots\}$ with $\#(K) = k, \#(N) = l$. We want to find some y adjacent to every vertex in K but not adjacent to any in L. Our first guess would be to form

$$s = \sum_{v_i \in K} 2^i$$

and take $y = v_s$. This y is certainly connected to all the K, and we never have $[s]_i = 1$ for $v_i \in L$, but if s is too small we may yet have $[i]_s = 1$ for some $v_i \in L$. Fortunately, this is easy to fix: just pick some $\ell > \max(K \cup L)$ and take

$$s' = s + 2^\ell.$$

This number will have the same small small bits as before, so $[s]_i$ will be 1 or 0 if $v_i \in K$ or L, respectively, and there will be no chance of $[i]_s = 1$ for $v_i \in L$, as $s \geq 2^\ell > \ell > \max(K \cup L) \geq \log_2 \max(L) + 1$ (the maximum number of binary digits of any element of L). □

The surprising thing is that the converse to this proposition is also essentially true. We need to restrict ourselves to countable models, but once that distinction is made, every model of EA is isomorphic to $\mathfrak{RG}$. This is a well-known model-theoretic property called *ω-categoricity.*

Proposition 2.8. *EA is ω-*categorical*. That is, if $\mathfrak{A}$ is countable and $\mathfrak{A} \models EA$, then $\mathfrak{A} \cong \mathfrak{RG}$.*

Proof. We use the power of the extension axioms to inductively build an isomorphism between any two countable models $\mathfrak{A}, \mathfrak{B} \models EA$. We will suppose without loss of generality that both $\mathfrak{A}$ and $\mathfrak{B}$ have the universe $\{0, 1, 2, \ldots\}$.

The process will work as follows:

- Begin with the trivial isomorphism i_0 from $\mathfrak{A}_0 = \emptyset$ to $\mathfrak{B}_0 = \emptyset$.

- On the k^{th} step of the process, for $k > 0$, perform one "$\mathfrak{A}$ to $\mathfrak{B}$" sub-step and one "$\mathfrak{B}$ to $\mathfrak{A}$" sub-step. The "$\mathfrak{A}$ to $\mathfrak{B}$" sub-step proceeds as follows: Find the smallest $a \in \mathfrak{A} \setminus \mathfrak{A}_{k-1}$ (that is, the smallest unmatched element of $\mathfrak{A}$). Let K denote the vertices of $\mathfrak{A}_{k-1}$ adjacent to a, and let L denote the vertices of $\mathfrak{A}_{k-1}$ non-adjacent to a. The extension axiom $EA_{\#(K),\#(L)}$, applied to the corresponding sets of vertices $i_{k-1}(K), i_{k-1}(L)$ in $\mathfrak{B}_{k-1} = i_{k-1}(\mathfrak{A}_{k-1})$, guarantees the existence of a vertex $b \in \mathfrak{B}$ such that when we extend i_{k-1} by sending a to b, we get an isomorphism i'_{k-1} from $\mathfrak{A}'_{k-1} = \mathfrak{A}_{k-1} \cup \{a\}$ to $\mathfrak{B}'_{k-1} = \mathfrak{b}_{k-1} \cup \{b\}$.

  To perform the "$\mathfrak{B}$ to $\mathfrak{A}$" substep, repeat these steps with the roles of $\mathfrak{A}$ and $\mathfrak{B}$ reversed, starting with $i'_{k-1} : \mathfrak{A}'_{k-1} \to \mathfrak{B}'_{k-1}$ and obtaining $i_k : \mathfrak{A}_k \to \mathfrak{B}_k$.

Since we kept on taking the smallest unmatched vertices from $\mathfrak{A}$ and $\mathfrak{B}$, every vertex will be picked eventually (by round k, we have $\{0, 1, \ldots, k\} \subseteq \mathfrak{A}_k$ and $\{0, 1, \ldots, k\} \subseteq \mathfrak{B}_k$). When we take the union of all the i_k, we will obtain a full isomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. □

To derive completeness from ω-categoricity, we need a classic result from model theory called the Löwenheim-Skolem theorem:

Theorem 2.9. *If a theory T has an infinite model, it has a countable model.*

Proposition 2.10. *EA is complete.*

Proof. Suppose there were some ϕ such that neither $EA \models \phi$ nor $EA \models \neg\phi$. Then both $EA \cup \{\phi\}$ and $EA \cup \{\neg\phi\}$ would be consistent and have models. Since every model of EA is infinite, this would imply by the Löwenheim-Skolem theorem that both $EA \cup \{\phi\}$ and $EA \cup \{\neg\phi\}$ would have countable models. Since EA is ω-categorical, so these countable models would both have to be isomorphic to $\mathfrak{RG}$. However, we cannot have both $\mathfrak{RG} \models \phi$ and $\mathfrak{RG} \models \neg\phi$. Thus, such a ϕ cannot exist, and we have proven that EA is complete. □

Now that we have EA complete, Proposition 2.3 applies, so we have proven Theorem 2.1, the zero-one law for first-order logic over graphs.

3 Fraïssé Limits and Random Structures

3.1 Homogeneousness

One way to look at the random graph is as a graph which contains all finite graphs in a particularly nice way. It is easy to see that $\mathfrak{RG}$ contains all finite graphs; the process of building up an arbitrary finite graph inside the random graph is essentially what we did in our proof of Proposition 2.8.

But this isn't enough to make the random graph ω-categorical. As an illustration of this, take the random graph and add on an isolated vertex to obtain a new graph $\mathfrak{RG} \cup \{v\}$. This new graph also contains every finite graph, but it is clearly not isomorphic to $\mathfrak{RG}$ (if nothing else, $EA_{1,0}$ doesn't work on the new vertex). The reason is that, in order to build up an isomorphism the way we did with extension axioms, we need to be able not only to embed any finite graph in our graph, but to be able to extend any embedding of a finite graph to an embedding of a larger finite graph. This sort of freedom follows from a powerful property known as homogeneousness, which happens to be the "particularly nice way" that the random graph contains all finite graphs.

Definition 3.1. *A structure* $\mathfrak{A}$ *is* homogeneous *if every isomorphism between finite substructures of* $\mathfrak{A}$ *extends to an automorphism of* $\mathfrak{A}$*. That is, if* $\mathfrak{B}, \mathfrak{C} \subseteq \mathfrak{A}$ *and there is an isomorphism* $i : \mathfrak{B} \to \mathfrak{C}$*, then there is an automorphism* $\hat{i} : \mathfrak{A} \to \mathfrak{A}$ *with* $\hat{i}|_{\mathfrak{B}} = i$*.*

$$\begin{array}{ccc} \mathfrak{B} & \xrightarrow{\subseteq} & \mathfrak{A} \\ {\scriptstyle \cong}\downarrow{\scriptstyle i} & & {\scriptstyle \cong}\downarrow{\scriptstyle \hat{i}} \\ \mathfrak{C} & \xrightarrow{\subseteq} & \mathfrak{A} \end{array}$$

Being homogeneous is enough to give the random graph ω-categoricity, and combining this with containing every finite graph is what allows us to use the random graph to prove the zero-one law for first-order logic over graphs. We can generalize this train of thought to prove the zero-one law for first-order logic over general structures, once we know how to build up an analogue to the random graph: a homogeneous random structure containing every finite structure. We do this using an even more general construction known as a Fraïssé limit.

3.2 Fraïssé Limits

Putting it as generally as possible, we are looking for a structure which contains a given class of structures as its class of finite substructures. However, this sort of construction won't be possible in general; only certain classes of structures appear as such a class. To give a name to such collections, we will call them *ages*.

Definition 3.2. *The* age *of a structure* $\mathfrak{A}$ *over a signature* σ *(finite, with no function symbols) is the collection* $\mathbf{K}$ *of all finite substructures of* $\mathfrak{A}$*. An* age*, in general, is any collection so obtained, up to isomorphism.*

We say "up to isomorphism" to mean that $\mathbf{K}$ is an age if there is some $\mathfrak{A}$ such that every finite substructure of $\mathfrak{A}$ is isomorphic to some element of $\mathbf{K}$, and visa versa.

We would like to find some set of properties of ages which taken together will imply that a certain collection is an age. A particularly natural pair can be found:

Definition 3.3. *A collection* $\mathbf{K}$ *of finite* σ*-structures satisfies*

- *the* hereditary property (HP) *if for every* $\mathfrak{A} \in \mathbf{K}$ *and every substructure* $\mathfrak{B} \subseteq \mathfrak{A}$*, there is some element of* $\mathbf{K}$ *isomorphic to* $\mathfrak{B}$*, and*

- *the* joint embedding property (JEP) *if, for every* $\mathfrak{A}, \mathfrak{B} \in \mathbf{K}$ *there exists some* $\mathfrak{C} \in \mathbf{K}$ *such that* $\mathfrak{A}$ *and* $\mathfrak{B}$ *are both embeddable in* $\mathfrak{C}$ *(isomorphic to substructures of* $\mathfrak{C}$*).*

Theorem 3.4. (Fraïssé's Weak Theorem) *A non-empty finite or countable collection* $\mathbf{K}$ *of finite* σ*-structures is an age iff it satisfies HP and JEP.*

Proof. ($\Rightarrow$) The hereditary property is true basically by transitivity, and the joint embedding property can be proven by taking two embedding in $\mathfrak{A}$ and taking their union.

($\Leftarrow$) Suppose $\mathbf{K} = \{\mathfrak{A}_0, \mathfrak{A}_1, \mathfrak{A}_2, \ldots\}$. We can construct a structure $\mathfrak{B}$ with age $\mathbf{K}$ by gluing together all the $\mathfrak{A}_i$. To be precise, we construct a tower of structures $\mathfrak{B}_i \in \mathbf{K}$ inductively as follows:

- $\mathfrak{B}_0 = \emptyset$.

- For $k \geq 0$, $\mathfrak{B}_{k+1}$ comes from jointly embedding $\mathfrak{B}_k$ and $\mathfrak{A}_k$.

We take $\mathfrak{B} = \bigcup_i \mathfrak{B}_i$. Since $\mathfrak{A}_k$ is embeddable in $\mathfrak{B}_{k+1}$ and $\mathfrak{B}_{k+1}$ is embeddable in $\mathfrak{B}$, every $\mathfrak{A}_i$ is a substructure of $\mathfrak{B}$. Conversely, any finite substructure of $\mathfrak{B}$ must be embeddable in some $\mathfrak{B}_i \in \mathbf{K}$, so by the hereditary property every finite substructure of $\mathfrak{B}$ is in $\mathbf{K}$. We have proven that $\mathbf{K}$ is an age. □

Every collection satisfying HP and JEP will be the age of some sort of structure. However, to make sure that the collection is the age of a unique homogeneous structure, we need a third property.

Definition 3.5. *A collection* $\mathbf{K}$ *of finite* σ*-structures satisfies the* amalgamation property (AP) *if for every* $\mathfrak{A}, \mathfrak{B}, \mathfrak{C} \in \mathbf{K}$ *with embeddings* $e : \mathfrak{A} \to \mathfrak{B}$ *and* $f : \mathfrak{A} \to \mathfrak{C}$ *there exists a* $\mathfrak{D} \in \mathbf{K}$ *with embeddings* $g : \mathfrak{B} \to \mathfrak{D}$ *and* $h : \mathfrak{C} \to \mathfrak{D}$ *such that* $ge = hf$*.*

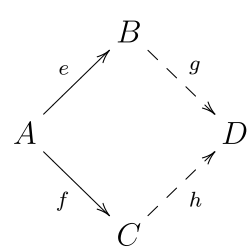

If every pair of structures in $\mathbf{K}$ have a shared embeddable substructure, JEP is a special case of AP, but this may not always be the case.

If an age $\mathbf{K}$ satisfies AP, it can be realized as the age of a homogeneous structure called a Fraïssé limit.

Theorem 3.6. (Fraïssé's Theorem) *Suppose* $\mathbf{K}$ *is a non-empty finite or countable collection of finite* σ*-structures which satisfies HP, JEP, and AP. Then there is a unique structure* $\mathfrak{D}$ *such that*

1. $\mathfrak{D}$ *is countable or finite,*
2. $\mathbf{K}$ *is the age of* $\mathfrak{D}$*, and*
3. $\mathfrak{D}$ *is homogeneous.*

Definition 3.7. *We call this the* Fraïssé limit *of* $\mathbf{K}$*, and denote it* $\lim \mathbf{K}$.

The proof of Fraïssé's Theorem is a bit involved, so it is omitted (a full proof can be found in section 7.1 of [1]).

3.3 The Random Structure

Let $\mathbf{K}$ be the class of all finite σ-structures. This $\mathbf{K}$ clearly satisfies HP, JEP, and AP, so it has a Fraïssé limit.

Definition 3.8. *The* random σ-structure $\textsc{Ran}(\sigma)$ *is the Fraïssé limit of the class* $\mathbf{K}$ *of all finite* σ*-structures.*

Just like we used graph extension axioms in Proposition 2.3 to prove the zero-one law for graphs, we will use a generalized sort of extension axiom to prove the zero-one law for general structures.

Given a σ-structure $\mathfrak{A}$ on the canonical $(n+1)$-element universe $\{0, \ldots, n\}$ we define the extension axiom $EA_{\mathfrak{A}}$. It will say that every n-element substructure of our model isomorphic to $\mathfrak{A}|_{\{0,\ldots,n-1\}}$ can be extended to a substructure isomorphic to all of $\mathfrak{A}$. Or, in first-order logic:

Definition 3.9. *The extension axiom* $EA_{\mathfrak{A}}$ *is the sentence*

$$\forall x_0 \cdots \forall x_{n-1}(\psi(x_0, \ldots, x_{n-1}) \Rightarrow \exists x_n(\chi(x_0, \ldots, x_{n-1}, x_n))),$$

where $\psi(x_0, \ldots, x_{n-1})$ *is a formula telling every relationship between the elements* $\{0, \ldots, n-1\}$ *and* $\chi(x_0, \ldots, x_{n-1}, x_n)$ *is a formula telling every relationship between the elements* $\{0, \ldots, n-1, n\}$.

The theory $EA[\sigma]$ *is the set of all these extension axioms:* $EA[\sigma] = \{EA_{\mathfrak{A}} \mid \mathfrak{A} \in \textsc{Struct}[\sigma]\}$.

When I say "formula telling every relationship between the elements $\{0, \ldots, n-1\}$, what I mean is a formula such that it will hold for $\{x_0, \ldots, x_{n-1}\}$ if and only if the induced substructure on that subset is isomorphic to $\mathfrak{A}|_{\{0,\ldots,n-1\}}$ by the obvious correspondence.

Proposition 3.10. *$\mu(EA_{\mathfrak{A}}) = 1$, for every $\mathfrak{A}$.*

Proof. A basic exercise in probability, essentially no different than the case for graphs. □

Proposition 3.11. *$EA[\sigma]$ is ω-categorical.*

Proof. It can easily be proven, using techniques similar to what we saw with the graph-based extension axioms, that any countable model $\mathfrak{A} \models EA[\sigma]$ must contain every finite σ-structure. This means that $\mathfrak{A}$ has the same age as $\textsc{Ran}(\sigma)$. In a very similar manner, we can prove that $\mathfrak{A}$ is homogenous. Thus, by Fraïssé's Theorem (3.6), $\mathfrak{A} \cong \textsc{Ran}(\sigma)$. □

Proposition 3.12. *$EA[\sigma]$ is complete.*

Proof. This follows from the ω-categoricity of $EA[\sigma]$ in exactly the same way as it did for EA in Proposition 2.10. □

Now that we have a complete theory of sentences which hold almost surely, we know that we have a zero-one law. The end-game proof here was practically identical to the case for graphs. The only real difference is that we used the powerful Fraïssé limit construction to construct an appropriately homogeneous model for our theory. This newfound power extends far beyond this one application. Scholars have used Fraïssé limits to prove all sorts of zero-one laws, as well as to reach into neighboring fields of order theory and group theory.

References

[1] Wilfred Hodges. *Model Theory.* Cambridge University Press, 1993.

[2] Leonid Libkin. *Elements of Finite Model Theory.* Springer, 2004.

[3] Joel Spencer. *Logic and Random Structures* (from *Finite Model Theory and Its Applications.* Springer, 2007.)