1 Random graph games

Let's play a game. Actually, let's play two different games. Both are about random graphs and their structural properties. As the author, I get to play the role of the house. Don't forget that the house always wins.

A **random graph** is, unsurprisingly, a randomly generated graph. Start with some vertex set V, and then for each pair of distinct vertices u, v, flip a coin to determine whether or not there is an edge connecting them. If the coin comes up heads, we get an undirected edge between u and v. Tails? No edge. There are more general notions of random graph that allow different edges to be generated with different probabilities, but we won't concern ourselves with them.

After flipping a coin for every edge, we have generated a simple, undirected graph with no self-loops or parallel edges. On the vertex set $V = \{1, \cdots, n\}$, each pair i, j may or may not have an edge between them, leading to $2^{\binom{n}{2}}$ possible different graphs that we could generate. But many of these graphs are essentially the same — the graphs are identical up to a renaming of the vertices. For example, the graph with a single edge between vertices 1 and 2 with no other edges is structurally identical to the graph with merely a single edge between 3 and 4. Try and count how many different graph structures there are on $\{1, \cdots, n\}$. You'll find that it's much more challenging than counting the number of graphs, but generating functions might help you out.

A **structural property** of a graph G is exactly a property that holds for all graphs with the same structure as G. For example, the properties of "having a 3-cycle", "being connected", and "having a clique of k vertices" are all examples of structural properties. On the other hand, any property that appeals to the names of the vertices, like "u and v are not joined by an edge", or "there is a 2-coloring in which u and v are the same color" are not structural.

We can now state our games.

The infinite game: You name a structural property of a graph, and the house declares whether or not a random graph on G with vertex set $\mathbb{N}$ will or will not have the property. A countable random graph is generated. If the house is correct, it wins.

The finite game: The finite game is identical to the infinite game, but this time the random graph is large, but finite. Say, $2^{2^{100}}$ vertices.

Will you play? And if so, what structural property will you pick? Analyzing these two games leads us down a wonderful garden path, seeing some highlights of graph theory, random geometry, first order model theory, probability, and one of our favorite theorems in all of mathematics.

2 The infinite game and the extension property

Let's fix some definitions and notations. A graph $G = (V, E)$ is a set of vertices V and a **edge relation** $E \subseteq V \times V$. We require the edge relation E to be symmetric (for all $u, v \in V$, we have $(u, v) \in E \implies (v, u) \in E$) and irreflexive (for all $v \in V$, we have $(v, v) \notin E$). Thus having $(u, v) \in E$ captures having an undirected non-self-loop edge between u and v. When $(u, v) \in E$, we will say that u and v are **adjacent**, and abbreviate this by $u \sim v$.

Given a graph $G = (V, E)$, a **subgraph** of G is a graph $G' = (V', E')$ such that $V' \subseteq V$ and $E' \subseteq E$. That is, G' is a graph with some (but not necessarily all!) vertices and edges from G. We write $G' \subseteq G$. There is a different, stronger notion of subgraph. We say that $G' = (V', E')$ is an **induced subgraph** of $G = (V, E)$ if $V' \subseteq V$ and $E' = E \cap (V' \times V')$. That is, G' is a subgraph of G with the additional property that for all $u, v \in V'$, if $u \sim v$ in G, then $u \sim v$ in G'. Alternatively, you can view an induced subgraph as the maximal subgraph on some selection of vertices.

Given two graphs $G = (V_G, E_G)$ and $H = (V_H, E_H)$, a **graph morphism** $f : H \to G$ is a function on the vertex sets $f : V_H \to V_G$ with the additional property that if $u \sim v$ in H, then $f(u) \sim f(v)$ in G. That is, adjacent edges in H are mapped to adjacent edges in H. Note that the converse need not be true, since $f : H \to G$ can send non-adjacent vertices of H to adjacent vertices of G. We say a graph morphism $f : H \to G$ is an **isomorphism** if $f : V_H \to V_G$ is a bijection, and has the property that

$$f(u) \sim f(v) \iff u \sim v$$

for all $u, v \in H$. That is, f simply *renames* the vertices of H, but leaves the structure of the graph unaltered. When G and H are isomorphic graphs, we write $G \cong H$. Finally, a **structural property** of a graph G is exactly a property that is true of all graphs G' such that $G' \cong G$.

Analyzing the infinite game

With these definitions and notations fixed, we can begin analyzing our games. We start with the infinite game. What does it take to win the infinite game? You'd need the house's guess about your structural property to be incorrect. Assuming the house is staffed by very talented mathematicians who know a lot about random graphs — and don't forget that I am the house — you'd maximize your chances by picking a structural property that holds very nearly half the time.

But finding such good structural property is a challenge. For example, let K_n denote the complete graph on n vertices — n vertices all connected to each other. Would "contains K_n as an induced subgraph" be a good structural property to bet on? Well, for any collection of n distinct vertices of a countable random graph, the probability that these vertices form K_n as an induced subgraph is exactly $2^{\frac{-n(n-1)}{2}}$. For large n, this is a very very small, but positive, probability. We can partition the vertices of a countable random graph into countably many disjoint sets of n vertices, each with a positive (and independent!) probability of forming K_n as an induced subgraph. Just as if you roll a $2^{\frac{n(n-1)}{2}}$ sided die infinitely many times you're guaranteed to eventually roll a 1, so too does it follow that a countable random graph is guaranteed to contain K_n as an induced subgraph. Ok, this is a slight misrepresentation. We should be saying something like "a countable random graph contains K_n as an induced subgraph with probability 1". But for the sake of clarity, we are happy to ignore such details. Anyway, "contains K_n as an induced subgraph" would not be a good structural property to bet on. The house declares "yes, the countable random graph will have K_n as an induced subgraph," and it will be correct every single time.

The previous argument easily generalizes to absolutely any finite graph. We have proved the following:

Proposition 1. *A countable random graph contains every finite graph as an induced subgraph.* □

The extension property

Unfortunately for you, we can extend this argument to essentially cover any "finite" structural property. This is because every countable random graph has a very special property. In a graph $G = (V, E)$, take disjoint sets of vertices $A, B \subseteq V$. Let $\mathrm{Ex}(A, B)$ denote the collection of vertices that are adjacent to every $a \in A$, and adjacent to no $b \in B$. That is,

$$\mathrm{Ex}(A, B) := \{v \in V \ : \ v \sim a \text{ for all } a \in A, \text{ and } v \not\sim b \text{ for all } b \in B\}.$$

G has the **extension property** if $\mathrm{Ex}(A, B)$ is non-empty for all finite disjoint sets $A, B \subseteq V$. We call a vertex $v \in \mathrm{Ex}(A, B)$ a **witness** to the extension property for A and B.

It's pretty straight forward to see that every countable random graph has the extension property. Fix the finite disjoint vertex sets A and B. Suppose that there are n and m vertices in A and B respectively. For any $w \notin A \cup B$, since each edge out of v is determined independently, w witnesses the extension property for A and B with probability $\frac{1}{2^{n+m}}$. As n and m grow big, this probability becomes vanishingly small — but it's always positive! And with *infinitely many* vertices outside of $A \cup B$ to check, we are again guaranteed to eventually find a vertex in $\mathrm{Ex}(A, B)$.

The extension property is of great import for the infinite game. Not only does it invalidate essentially every finite structural property as a good bet. It even has huge implications for the wholistic structrue of countable random graphs. To start getting comfortable with the extension property, we prove two quick propositions.

Proposition 2. *Let G be a graph with the extension property. Then G is empty or has infinitely vertices.*

This is baked in to the extension property. Given distinct vertices $B := \{v_1, \cdots, v_n\}$ of G, there is a vertex $w \in \mathrm{Ex}(\varnothing, B)$ that is not connected to any v_j. Thus if G has $n > 0$ vertices, it must have $n + 1$ vertices. It follows that if G has any vertices at all, it must have infinitely many.

Proposition 3. *Let G be a graph with the extension property. For disjoint finite sets of vertices A, B, the set of witnesses $Ex(A, B)$ is infinite.*

The extension property tells us that $\mathrm{Ex}(A, B)$ is non-empty. So pick a witness $w_1 \in \mathrm{Ex}(A, B)$. Now consider the set $\mathrm{Ex}(A, B \cup \{w_1\})$. The extension property again ensures that this set is non-empty, so we can select w_2. Further, $\mathrm{Ex}(A, B \cup \{w_1\})$ is a subset of $\mathrm{Ex}(A, B)$, so $w_2 \in \mathrm{Ex}(A, B)$. Iterating this construction and selecting $w_{n+1} \in \mathrm{Ex}(A, B \cup \{w_1, \cdots, w_n\})$ builds an infinite set of witnesses $\{w_1, w_2, \cdots\} \subseteq \mathrm{Ex}(A, B)$. This proves the proposition.

We already showed that a countable random graph contains every finite graph as an induced subgraph. Using the extension property, we can prove something much much stronger.

Theorem 4. *Suppose that G has the extension property, and H is a finite or countable graph. Then H is an induced subgraph of G.*

PROOF: Let H be a countable graph. Denote the vertices of H by $\{h_0, h_1, h_2, \cdots\}$. We can inductively pick vertices $\{g_0, g_1, g_2, \cdots\}$ of G such that $g_i \sim g_j$ if and only if $h_i \sim h_j$ for all $i, j \in \mathbb{N}$. We start by picking any vertex g_0 of G. Now suppose we've already picked vertices $g_0, \cdots, g_{n-1}$ of G such that $g_i \sim g_j$ if and only if $h_i \sim h_j$ for all $0 \leq i, j \leq n - 1$. Set $A_n, B_n \subseteq \{g_0, g_1, \cdots, g_{n-1}\}$ by:

$$g_i \in A_n \iff h_i \sim h_n$$

$$g_j \in B_n \iff h_j \not\sim h_n$$

Now we pick g_n to be a witness from $\mathrm{Ex}(A_n, B_n)$. By construction, for $i \le n-1$ we have that $g_n \sim g_i$ if and only if $h_n \sim h_i$. By induction, we build a set of vertices $\{g_0, g_1, g_2, \cdots\}$ such that $g_i \sim g_j$ if and only if $h_i \sim h_j$. Thus we see that H is an induced subgraph of G. □

So any structural property we want to bet on can't be a claim about an induced subgraph. Things are starting to seem hopeless, but there are other properties we can investigate. We haven't looked at colorability, for example, or more exotic topological properties like homology and homotopy groups. Before we look at other properties, let's pin down a few concrete (read: not random) graphs with the extension property.

Graphs with the extension property

It's not very hard to find graphs with the extension property. We give you three.

Example 5. *Let G be the graph with vertex set $\mathbb{N}$, and an edge between i and j if the j'th bit of the binary representation of i is 1, or vice versa. G has the extension property.*

To prove this, given disjoint sets of natural numbers $A = \{a_1, \cdots, a_n\}$ and $B = \{b_1, \cdots, b_m\}$, we must construct a witness to the extension property. To do this, first pick $k = \max\{a_1, \cdots, a_n, b_1, \cdots, b_m\}$. Now take

$$w := 2^{k+1} + 2^{a_1} + \cdots + 2^{a_n}.$$

The a_j'th binary digit of w is 1 by construction, so $w \sim a_j$ for each $j = 1, \cdots, n$. On the other hand, the b_j'th binary digit of w is 0 for all $j = 1, \cdots, m$. Additionally, for all j, we have $w > \log_2(b_j)$, so the w'th binary digit of b_j is 0. Thus $w \not\sim b_j$. This proves $w \in \mathrm{Ex}(A, B)$, so G has the extension property.

Example 6. *Let G be the graph with vertex set $V := \{p \in \mathbb{N} : p \text{ is prime and } p = 1 \mod 4\}$, and an edge between p and q if and only if p has a square root modulo q. G has the extension property.*

It might not be clear that this edge relation is symmetric. This is a consequence of the **law of quadratic reciprocity**.

Theorem 7 (law of quadratic reciprocity)**.** *Let p and q be odd primes. The following holds:*

$$\big((p = 1 \mod 4) \text{ or } (q = 1 \mod 4)\big) \iff \big((p \text{ is a square modulo } q) \iff (q \text{ is a square modulo } p)\big).$$

Since p and q are both primes congruent to 1 modulo 4, our edge relation is symmetric.

The law of quadratic reciprocity is a famously hard problem. It was a conjecture of Euler, but not a theorem of Euler! It took the genius of Gauss to prove the result. It's a bit too much of a detour away from random graphs to prove it here, but there is a beautiful (and concise!) proof by Veklych [**?**].

Now that we know the edge relation is symmetric, proving that G has the extension property is easy. Given disjoint vertex sets $A = \{a_1, \cdots, a_n\}$ and $B = \{b_1, \cdots, b_m\}$, we construct a witness to the extension property. Start by picking values $s_1, \cdots, s_n, t_1, \cdots, t_m$ such that s_i is a perfect square modulo a_i, and t_j is not a perfect square mod b_j. By **Sunzi's theorem**, there is a solution v to the following system of equations:

$$x = 1 \mod 4$$

$$
\begin{aligned}
x &= s_i \mod a_i \qquad \text{for } i = 1, \cdots, n \\
x &= t_j \mod b_j \qquad \text{for } j = 1, \cdots, m.
\end{aligned}
$$

if v is prime, then we're done! Otherwise, we use **Dirichlet's theorem on arithmetic progressions** to assert that there is some prime p of the form $p = v + k(4a_1 \cdots a_n b_1 \cdots b_m)$ for some $k \in \mathbb{N}$. The prime p is a witness to the extension property.

Example 8. *The hereditary finite sets with the symmetrized membership relation "ϵ" has the extension property.*

Let's unpack that. We start by inductively defining the a collection of sets V_n for each $n \in \mathbb{N}$ by:

$$
\begin{aligned}
V_0 &:= \varnothing; \\
V_{n+1} &:= \mathcal{P}(V_n); \\
V_\infty &:= \bigcup_{n=0}^{\infty} V_n
\end{aligned}
$$

where $\mathcal{P}$ denotes the powerset operation. The collection V_∞ is known as the **hereditary finite sets**. These are all the finite sets of finite depth that can be built from the empty set and the powerset operation. We take the hereditary finite sets to be the vertices of our graph G, with an edge joining x and y if $x \in y$ or if $y \in x$.

Given disjoint vertex sets $A = \{a_1, \cdots, a_n\}$ and $B = \{b_1, \cdots, b_m\}$, to construct a witness to the extension property, pick any hereditary finite set x such that $x \notin B$ and for all $z \in b_1 \cup \cdots \cup b_m$, we have $x \notin z$. Now consider the set $w = \{x, a_1, \cdots a_n\}$. This is a hereditary finite set with $a_i \in w$ for all $i = 1, \cdots, n$. Further, we have by hypothesis that $b_j \notin w$ for all $j = 1, \cdots, m$. Finally, if $w \in b_j$, then $x \in w \in b_j$, contradicting our construction of x. Thus w is a witness to the extension property.

3 The Erdős-Rényi-Rado graph

Based on the previous section, it looks like examples of countable graphs with the extension property abound. The countable random graph has the extension property, as do the three countable graphs we just analyzed. But we've been hiding something from you — these are all the same graph up to renaming the vertices.

Theorem 9. *Let G and G' be countable graphs with the extension property. Then G and G' are isomorphic graphs.*

Before we prove this theorem, let's give some commentary. The unique countable graph with the extension property is known as the **Erdős-Rényi-Rado graph**, or just the Rado graph for short. In the last section, we gave three different constructions of the Rado graph. The countable random graph is also the Rado graph. Note what this means for the infinite random graph game. The flipping of the coin doesn't matter; The Rado graph will be generated, and the house will win every time. We warned you! All structural properties are predetermined, and there are no good properties to bet on. Also, take note that our proof that the countable random graph has the extension property didn't actually rely the fact that the coin was fair. If the coin comes up heads with any probability $p \in (0, 1)$, then the dealer generates the random graph. Even though the graph generated by a $p = 2^{-100}$ coin *feels like* it should be much more sparse than the graph generated by the $p = .999$ coin, it's the same graph!

Take a moment to think about how weird this is. As n gets larger, the number of different graphs on n vertices, up to isomorphism, grows exponentially in n. We get more and more graphs until we actually *reach infinity* and the number of possibilities collapses to one. It's really quite shocking the first time you see it.

In some circles, this property of only having one countable structure up to isomorphism with a given property is known as ω-**categoricity**. Here, ω is just the ordinal notation for the size of the natural numbers. Another example of an ω-categorical structure is $\mathbb{Q}$, viewed as an ordered set. If $(L, \leq)$ is any other countable dense ordered set without endpoints, then there is an order preserving bijection between L and the rational numbers $\mathbb{Q}$. These kinds of properties are quite rare to come across in the wild, but they're quite valuable. ω-categoricity results are basically all proved the same way — by the **back-and-forth argument**. This is a proof technique that you should be aware of. It doesn't come up a whole lot, but sometimes it's the silver bullet you need.

The back-and-forth argument starts with a pair of countable structures X and Y, and an isomorphism $f_0 : X_0 \to Y_0$ for finite substructures $X_0 \subseteq X$ and $Y_0 \subseteq Y$. Right now, by isomorphism we simply mean that f_0 is a structure-preserving bijection on these subsets. Now pick any $x \in X \setminus X_0$ and $y \in Y \setminus Y_0$. We must show that there are substructures $X_0 \cup \{x\} \subseteq X_1 \subseteq X$ and $Y_0 \cup \{y\} \subseteq Y_1 \subseteq Y$ and an isomorphism $f_1 : X_1 \to Y_1$ that extends f_0. That is, we can add any point to the domain and any point to the codomain, and extend our partial isomorphism so as to include those points in the domain and range. If you can prove this, it immediately follows that X and Y must be isomorphic. Just keep going back and forth, extending the domain and range! After infinitely many steps (and secretly a transfinite induction), you will have built an isomorphism $f : X \to Y$. Let's do it explicitly for the Rado graph.

Proof of theorem 9: Suppose G and H are countable graphs with the extension property. Let $G_0 \subseteq G$ and $H_0 \subseteq H$ be finite induced subgraphs of G and H, and let $f_0 : G_0 \to H_0$ be an isomorphism. Let u be a vertex in $V(G) \setminus V(G_0)$ and v' be a vertex in $V(H) \setminus V(H_0)$. We must pick an image for u and a preimage for v'.

To pick an image for u, we use the extension property. Partition the vertices of $V(G_0) = A \cup B$, where $u \sim a$ for every $a \in A$, and $u \not\sim b$ for every $b \in B$. Since f_0 is an isomorphism, we may partition the vertices of H_0 into $f_0(A) \cup f_0(B)$. By the extension property, we may pick a vertex u' of H which is joined to every vertex in $f_0(A)$ and to no vertex in $f_0(B)$. We select u' to be the image for u.

To pick a preimage for v', we again use the extension property. Partition the vertices of $H_0 \cup \{u'\}$ into $C \cup D$, where v' is adjacent to every vertex in C and to no vertex in D. Since f_0 is an isomorphism, $f_0^{-1}(C) \cup f_0^{-1}(D)$ partitions the vertices of $G_0 \cup \{u\}$. By the extension property, we may pick a vertex v of G adjacent to every vertex in $f_0^{-1}(C)$ and to no vertex in $f_0^{-1}(D)$. We select v to be the preimage of v'.

Let G_1 be the induced subgraph of G with vertices $V(G_1) = V(G_0) \cup \{u, v\}$. Similarly, let H_1 be the induced subgraph of H with vertices $V(H_1) = V(H_0) \cup \{u', v'\}$. Now extend our map f_0 to $f_1 : G_1 \cup \to H_1$ defined by:

$$f_1(u) = u'$$
$$f_1(v) = v'$$

$$f_1\big|_{G_0} = f_0$$

By the way we chose u' and v, it is clear that f_1 is an isomorphism. $\square$

Other infinite random graph games?

So the infinite game is a bust. We've revealed the ace up our sleave, and you'll never play against us. But what if we changed the game to use some other infinite set of vertices? For example, what if we generated the random graph on $\mathbb{R}$ instead of $\mathbb{N}$?

Don't get confused about what Theorem 9 is saying. The Rado graph is not the only infinite graph with the extension property. There are graphs with λ-many vertices that have the extension property for any infinite cardinality λ. The crux of the argument is that given any graph $G_0 = (V_0, E_0)$, we can build a graph $\overline{G}$ that satisfies the extension property by adding $|V_0|$-many vertices to G_0. To do this, we build a sequence of subgraphs $G_0 \subseteq G_1 \subseteq G_2 \subseteq \cdots$. The graph G_{n+1} is constructed from $G_n := (V_n, E_n)$ by adding a vertex v_X for every finite set of vertices $X \subseteq V_n$. The edges of E_{n+1} has every edge from E_n, plus edges connecting v_X to every vertex $x \in X$. Finally, we set $\overline{G} := \bigcup_{n\in\mathbb{N}} G_n$ to be the limit of these graphs, and call it the **extension closure** of G_0. To see that $|\overline{V}| = |V_0|$, note that we add $|V_0|$-many vertices at every stage of the construction. There are countably many stages, so the vertex set of $\overline{G}$ has the same cardinality as the vertex set of G_0. To see that $\overline{G}$ has the extension property, take disjoint finite sets of vertices A, B of $\overline{G}$. Since A and B are finite, there is some $n \in \mathbb{N}$ such that every vertex in $A \cup B$ is in G_n. By construction, there is a vertex $v_A \in G_{n+1} \subseteq \overline{G}$ such that v_A is adjacent to every $a \in A$, but to none of the vertices in B. Thus $v_A \in \mathrm{Ex}(A, B)$.

The extension closure can further be used to show that for every cardinality $\lambda > |\mathbb{N}|$, there are graphs G, G' with λ-many vertices, both with the extension property, such that $G \not\cong G'$. For example, let I_λ be the edgeless graph with λ-many vertices, and let K_λ be the complete graph with λ-many vertices. Consider the extension closures $\overline{I_\lambda}$ and $\overline{K_\lambda}$. These graphs cannot be isomorphic, because $\overline{I_\lambda}$ does not contain K_λ as an induced subgraph. Let W be any collection of λ-many vertices from $\overline{I_\lambda}$. We claim that there are vertices $u, v \in W$ that are not adjacent. To see this, note that there must be a pair of vertices $u, v \in W$ that were added at the same stage of the construction of the extension closure of I_λ. By construction, these vertices cannot be adjacent. Since W was arbitrary, the graph $\overline{I_\lambda}$ cannot contain K_λ as an induced subgraph. However, $\overline{K_\lambda}$ clearly does contain K_λ as an induced subgraph, so $\overline{I_\lambda} \not\cong \overline{K_\lambda}$.

The two preceding paragraphs demonstrate that the magic of the Rado graph only happens for countable graphs. If you play the random graph game on an uncountable vertex set, the house is not guaranteed to generate the same graph every time. So, game on? Unfortunately, no. Game not on. Game still off. Talking about such uncountable random graphs is probabilistically problematic. In generating an uncountable random graph, we're taking the intersection of uncountably many independent events that determine the adjacency structure. The modern framework of probability is able to elegantly handle countable intersections, but things get weird with uncountable intersections. Part of this weirdness is structural, and baked into the Kolmogorov axioms and the definition of "probability itself". The other part of this weirdness is purely arithmetic — an uncountable product of numbers strictly between zero and one is always zero, giving no nontrivial probabilities. But if we somehow found a way to make probabilistic sense of an uncountable random graph, I'll be waiting for you at the felt-top table between the poker room and the penny slots.

4 Properties of the Rado graph

Before we analyze the finite game, we take a detour into the properties of the Rado graph. It's a really cool graph, with some mind boggling properties. And if you open your own casino where *you get to be the house*, you'll need to know what to say when some patsy comes in and starts rattling off structural properties.

At this point we know that the Rado graph $R = (V_R, E_R)$ is the unique countable graph with the extension property, and contains every countable subgraph as an induced subgraph. But what does it look like? What properties does it have? While we gave a few different explicit descriptions of R, these characterizations are pretty opaque. We certainly can't tell when two primes p and q are quadratic residues with residues with respect to each other just by looking. But we can use the extension property to prove a variety of characterizing properties of R. For example, a graph is **connected** if you can walk from any vertex to any other along edges of the graph. Is R connected? Well of course it is. The vertices u and v are connected by the path through any $w \in \text{Ex}(\{u, v\}, \varnothing)$.

Cliques

A **clique** in R is a collection of vertices $C \subseteq V_R$ such that every pair of vertices in C are adjacent. Call a clique C **maximal** if for any other clique C' such that $C \subseteq C'$, we have $C = C'$. Let v be a vertex of R, and let C_v be a maximal clique containing v. This maximal clique isn't unique, but it does exist. By the extension property, we know that C_v is infinite. Since there are infinitely many vertices that aren't adjacent to v, it follows that there are infinitely many disjoint infinite cliques in R. A similar argument shows that there are infinitely many disjoint infinite independent sets of vertices.

Colorings

The extension property also can be used to show that the Rado graph is hard to color. An **n-coloring** of a graph $G = (V, E)$ is a function $c : V \to \{1, \cdots, n\}$ such that if $v \sim w$, we have $c(v) \neq c(w)$. We think of the codomain $\{1, \cdots, n\}$ as a set of n different colors. The function c assigns a color to each vertex such that no two adjacent vertices share the same color. By the extension property, R is not n-colorable for any $n \in \mathbb{N}$. To see this, suppose that $c : V_R \to \{1, \cdots, n\}$ is an n-coloring. Without loss of generality, suppose that c is surjective. Pick vertices $v_1, \cdots, v_n$ such that $c(v_i) = i$. By the extension property, there is some vertex $w \in \text{Ex}(\{v_1, \cdots, v_n\}, \varnothing)$. What color can be assigned to w? We would need an $(n + 1)$'th color for w — a contradiction.

On the other hand, the Rado graph is trivially ω-colorable. Simply assign each vertex it's own color.

Self-similarity

The Rado graph is highly self-similar. In particular, it is isomorphic to many different subgraphs and supergraphs of itself. A **finite vertex pruning** of R is a subgraph $R' \subseteq R$ obtained by deleting finitely many vertices of R, and all the edges connected to the deleted vertices. Given such a vertex pruning $R' \subseteq R$, it is clear that R' still has the extension property. It follows that $R' \cong R$. Similarly, adding or deleting finitely many edges in the Rado graph will give you another isomorphic copy of the Rado graph. The strongest result in this vein is that the Rado graph has the pigeonhole property. A partition of a graph G is a pair of induced subgraphs, covering all the vertices of G, with no edges connected the induced subgraphs. That is, a partition of $G = (V, E)$ is a pair of induced subgraphs $G_1 = (V_1, E\big|_{V_1})$ and $G_2 = (V_2, E\big|_{V_2})$ such that $V_1 \cup V_2 = V$ and $V_1 \cap V_2 = \varnothing$. A graph G is said to have the **pigeonhole property** if for any partition (G_1, G_2), either $G_1 \cong G$ or $G_2 \cong G$. For example, the complete graph on $\mathbb{N}$ (denoted $K_\mathbb{N}$) and the edgeless graph on $\mathbb{N}$ (denoted $I_\mathbb{N}$) both have the pigeonhole property.

Proposition 10. *The Rado graph has the pigeonhole property.*

Proof: We proceed by contradiction. Suppose (R_1, R_2) is a partition of the Rado graph R such that neither R_1 nor R_2 is isomorphic to R. It follows that neither R_1 nor R_2 can have the extension property. We may pick finite disjoint sets of vertices $A_1, B_1 \subseteq V_1$ and $A_2, B_2 \subseteq V_2$ such that $\mathrm{Ex}_{R_1}(A_1, B_1)$ and $\mathrm{Ex}_{R_2}(A_2, B_2)$ are both empty. Here, the subscripts on Ex are keeping track of which graph we're working in. Since R has the extension property, we may find a witness:

$$w \in \mathrm{Ex}_R(A_1 \cup A_2, B_1 \cup B_2).$$

Now which subgraph is the vertex w in? If $w \in V_1$, then w would necessarily be a witness in $\mathrm{Ex}_{R_1}(A_1, B_1)$. Similarly, if $w \in V_2$, then w would necessarily be a witness in $\mathrm{Ex}_{R_2}(A_2, B_2)$. In either case, we have a contradiction. □

Perhaps surprisingly, these three graphs — the countable clique $K_\mathbb{N}$, the countable edgeless graph $I_\mathbb{N}$, and the Rado graph R — are the only countable graphs with the pigeonhole property up to isomorphism. To see this, let $G = (V_G, E_G)$ be a countable graph with the pigeonhole property. Call a vertex $v \in V$ **isolated** if it is adjacent to no other vertex, and **universal** if it is adjacent to every other vertex. Partition G into (I, L, U), where I is the induced subgraph of all the isolated vertices of G, U is the induced subgraph of all the universal vertices of G, and L is the induced subgraph on every vertex that is neither universal nor isolated. By the pigeonhole property, exactly one of these graphs is isomorphic to G, and the other two are empty. In particular, if $G \cong I$, then every vertex of G is isolated, and $U = L = \varnothing$. Similarly, if $G \cong U$, then every vertex of G is universal, and $I = L = \varnothing$. Finally, if $G \cong L$, then G has no universal vertices and no isolated vertices, and $I = U = \varnothing$.

Now suppose that G fails to have the extension property. Let A, B be disjoint finite sets of vertices such that $\mathrm{Ex}_G(A, B) = \varnothing$. Suppose we partition $A = A_1 \cup A_2$ and $B = B_1 \cup B_2$, allowing the possibility of $A_i = \varnothing$ or $B_j = \varnothing$. Let G_1 and G_2 be the induced subgraphs of G with vertex sets

$$V_1 := A_1 \cup B_1 \cup (V_G \setminus (\mathrm{Ex}(A_1, B_1) \cup A_2 \cup B_2))$$
$$V_2 := A_2 \cup B_2 \cup (V_G \setminus (\mathrm{Ex}(A_2, B_2) \cup V_1).$$

By design, G_1 and G_2 form a partition of G, so one of these graphs is isomorphic to G by the pigeonhole property. Since the partition was arbitrary, we can pick a partition where one of A, B are empty, and the other is a singleton. That is, there is a vertex $v \in V_G$ for which $\mathrm{Ex}(v, \varnothing) = \varnothing$ or $\mathrm{Ex}(\varnothing, v) = \varnothing$. In the first, case, v is an isolated vertex, making $G \cong I_\mathbb{N}$. In the second case, v is a non-isolated vertex, making $G \cong K_\mathbb{N}$.

There are many other wonderful properties of the Rado graph that we can't get into here. But we will make some suggestions for further reading at the end of the chapter.

5 The finite game and first order properties of graphs

We now shift our analysis toward the finite game. This time, we do have a game on our hands. For example, consider the property "has at least k edges" as a structural property. For $k = 1$, our enormous finite random graph with $2^{2^{100}}$ vertices could fail to have any edges, but the probability is negligible. Similarly, for $k = (2^{2^{100}}$choose 2), our enormous random graph *theoretically* could have every possible edge, but in the same way that we *theoretically* could quantum tunnel our way to China. In the middle, at $k = \frac{1}{2}(2^{2^{100}}$choose 2), the probability is exactly 1/2.

This argument, while accurate to the way we presented the finite random graph way back in the introduction, goes against the spirit of the game. True gamblers know that gambling isn't about the money, but sportsmanship. The issue is that "has at least $\frac{1}{2}(2^{2^{100}}$choose 2)" being a good bet is an artifact of our specific choice for the number of vertices. If you made this bet and the house said "oops, we're sorry. We meant $2^{2^{2^{100}}}$ vertices, not $2^{2^{100}}$," you would have made a very bad choice for your structural property. The point of the finite random graph game is to bet on the behavior of arbitrarily large (but still finite!) random graphs. Not on random graphs of a specific size.

On the other hand, consider the property "has an even number of edges". In this case, for any large N, the probability that a random graph on N vertices has an even number of edges is exactly 1/2. This property doesn't violate the spirit of the game, and would make for a good bet.

How do we distinguish good, honorable bets from implementation specific nonsense? You'd like to put a uniform distribution over large numbers, but of course there is no such distribution. Instead, we can leverage the notion of **density**. For a given property, let μ_N denote the probability that a random graph on N vertices has the property. We're interested in the limit $\lim_{N\to\infty} \mu_N$. This essentially captures the probability that an arbitrarily large graph has the desired property. Looking back, we see that the properties "has at least $\frac{1}{2}\binom{k}{2}$ edges" and "has an even number of edges" have densities of 0 and 1/2 respectively.

So are we done? We found a good bet for the finite random graph game. Fortunately, we're only just beginning our analysis of the finite game. What's interesting about the finite game isn't that there is a winning bet, but which bets could plausibly be winners. For example, what if we restrict ourselves to so-called "first order" properties of graphs? While most structural properties of graphs are not first order, the first order properties are well behaved in ways that allow us to prove a miraculous result called the zero-one law for finite graphs. Every structural property carves the collection of finite graphs into two pieces: the graphs that have the property, and the graphs that don't. The zero-one law says that for first order properties, this partition is boring — one piece will contain almost every graph in the sense of density. In order to prove this incredible theorem, we'll need to do a quick tour through some foundational results in mathematical logic.

First order properties

So what exactly is a first order property? A **first order property** ϕ, also known as a first order sentence, is a true/false statement about a graph built up from finitely many of the following pieces:

1. universal quantification over vertices $\forall v$ (read "for all vertices v");

2. existential quantification over vertices $\exists v$ (read "there exists a vertex v");
3. the edge relation $x \sim y$ (read "vertex x is adjacent to vertex y");
4. conjunction, $\wedge$ (read "and");
5. disjunction, $\vee$ (read "or");
6. implication, $\implies$ (read "implies");
7. equality, $=$ (read "equals");
8. negation $\neg$ (read "not").

We also highlight what you can't use in a first order property. You can't quantify over edges directly — only over vertices. That is, there's no way to directly say "there exists an edge". However, this can be circumvented by cleverly quantifying over vertices and using the edge relation. For example, the property $\phi := \exists x \exists y (x \sim y)$ (read "there are vertices x and y such that x is adjacent to y") is semantically equivalent to having an edge. The other important thing you can't do is quantify over *sets* of vertices. That is, you can't say that there is a set of vertices with some property, or that for all sets of vertices, some property holds. Some example will make things more clear.

Example 11. *The following are first order properties of graphs, with there formal characterizations:*

- *There are two distinct vertices:*
$$\exists x \exists y (x \neq y)$$
(read "there are vertices x and y such that $x \neq y$");
- *Every vertex is adjacent to exactly one vertex:*
$$\forall x \exists y ((x \sim y) \wedge (\forall z ((x \sim z) \implies (z = y))))$$
(read "for all vertices x, there is a vertex y such that $x \sim y$, and if z is any other vertex adjacent to x, then $z = y$");
- *The four cycle is a subgraph:*
$$\exists a \exists b \exists c \exists d ((a \sim b) \wedge (b \sim c) \wedge (c \sim d) \wedge (d \sim a))$$
(read "there are vertices a, b, c, d such that $a \sim b$, $b \sim c$, $c \sim d$, and $d \sim a$);
- *The four cycle is an induced subgraph:*
$$\exists a \exists b \exists c \exists d ((a \sim b) \wedge (b \sim c) \wedge (c \sim d) \wedge (d \sim a) \wedge \neg((a \sim c) \vee (b \sim d)))$$
(read "there are vertices a, b, c, d such that $a \sim b$, $b \sim c$, $c \sim d$, and $d \sim a$), and neither $a \sim c$ or $b \sim d$);
- *The four cycle is not a subgraph:*
$$\neg \exists a \exists b \exists c \exists d ((a \sim b) \wedge (b \sim c) \wedge (c \sim d) \wedge (d \sim a))$$
(read "there are not vertices a, b, c, d such that $a \sim b$, $b \sim c$, $c \sim d$, and $d \sim a$);

- *More generally, the presence (or non-presence) of any finite graph as an (induced) subgraph is a first order property by similarly specifying vertices and adjacencies;*

- *There are exactly n vertices:*

$$\exists v_1 \exists v_2 \cdots \exists v_n \left(\left(\bigwedge_{i \neq j} \neg(v_i = v_j) \right) \wedge \forall w \left(\bigvee_{i=1}^{n} (w = v_i) \right) \right)$$

(read "there are vertices $v_1, \cdots v_n$ such that $v_i \neq v_j$ for all $i \neq j$, and every vertex w is equal to one of the v_i");

- *Fix $n, m \in \mathbb{N}$. For all disjoint finite sets of vertices $A = \{a_1, \cdots, a_n\}$ and $B = \{b_1, \cdots, b_m\}$, there is a witness $w \in Ex(A, B)$:*

$$E_{n,m} := \forall a_1 \cdots \forall a_n \forall b_1 \cdots \forall b_m \Bigg(\Bigg(\left(\bigwedge_{i,j} \neg(a_i = b_j) \right) \wedge \left(\bigwedge_{i \neq j} \neg(a_i = a_j) \right) \wedge \left(\bigwedge_{i \neq j} \neg(b_i = b_j) \right) \Bigg) \implies \exists w \Bigg(\left(\bigwedge_{i=1}^{n} w \sim a_i \right) \wedge \left(\bigwedge_{j=1}^{m} \neg(w \sim b_j) \right) \Bigg) \Bigg).$$

This sentence is very long, and we encourage you to try to understand why it encodes the extension property for specific sizes of A and B. This first order property is important enough that we give it a name: $E_{n,m}$ is the ***(n,m)-extension property****.*

Example 12. *The following are not first order properties:*

- *While having any finite graph as an (induced) subgraph is a first-order property, the same is not true for infinite subgraphs. For example, having an infinite clique is not a first order property. To try and state this in the restricted language of first order properties we'd need to use infinitely many existential quantifiers.*

- *Having finitely many vertices is not a first order property. While having exactly n vertices is a first order property, there is no way to pick out "finiteness" in a single sentence. Let F_n be the first order property we defined above that corresponds to having exactly n vertices. You might try and define finiteness as a first order property by "$\exists n(F_n)$," but we can't do this! The natural number n isn't a vertex, so we can't apply an existential quantifier to it. What about "$F_0 \vee F_1 \vee F_2 \vee \cdots$"? Again, this isn't a valid first order property, as every first order sentence has finite length.*

- *Since a graph has infinitely many vertices if and only if it doesn't have finitely many vertices, the property of having infinitely many vertices is not first order.*

- *Similarly, even though the (n, m)-extension property is a first order property for all $n, m \in \mathbb{N}$, the extension property is not. We can't do universal quantification over n and m and make the extension property a first order property. However, we can say that a graph G has the extension property if and only if it has the (n, m)-extension property for all pairs $n, m \in \mathbb{N}$. Thus having the extension property is equivalent to satisfying an infinite family of first order properties.*

- *Being 2-colorable (or n-colorable more generally) is not a first order property. To define a two coloring as a first order property, you might try to show the graph is* ***bipartite****. That is, the vertices can be partitioned into two disjoint sets X and Y such that no two vertices in X are adjacent, and no two vertices in Y are adjacent. But this kind of reasoning can't be accommodated in the first order world! We can't quantify over subsets of vertices — only over single vertices.*

- *Being connected is not a first order property. Say that two vertices u and v in a graph G are d-far apart if the shortest path from u to v uses d edges of G. So every vertex is 0-far from itself, and adjacent vertices are 1-far apart. While "all vertices are no more than d-far apart" is a first order property, there is no way to assemble these into a first order sentence that exactly captures connectedness.*

So as you can see, first order properties are pretty restrictive. Most "natural" structural properties of graphs are not first order. But what we lose in expressiveness we gain ten-fold in behavior. In the coming sections, we will frequently rely on the logical properties of first order properties to prove incredible results about random graphs.

Going forward, when discussing structural properties of graphs, we will generally reserve lower-case Greek characters for first order properties, while we use upper case Greek characters for arbitrary (not necessarily first order) properties.

Elementary equivalence

Given a graph G and a first order property ϕ, the property ϕ is either true or not true of G. If ϕ is true of G, we say G **satisfies** ϕ, and write $G \vDash \phi$. We reserve this vocabulary and notation for first order properties. Since "not" is in our first order vocabulary, we see that G does not satisfy ϕ if and only if G satisfies $\neg\phi$. Often times, it is convenient to only consider graphs up to their first-order properties. We say that two graphs G and G' are **elementarily equivalent**, written $G \simeq G'$, if they satisfy exactly the same first order properties.

Clearly if G and G' are isomorphic, they are elementarily equivalent, but is the converse true? Or can non-ismomorphic graphs be elementarily equivalent? Let G and G' be elementarily equivalent graphs. First, notice that since "has n vertices" is a first order property, either G and G' are both finite with the same number of vertices, or they're both infinite. Since every graph is a subgraph of itself, and "contains the finite graph H as a subgraph" is a first order property, if G and G' are finite graphs, they both must be subgraphs of each other. Consequently, for finite graphs, $G' \simeq G \implies G \cong G'$. On the other hand, things get a little more weird when G and G' are infinite graphs.

Theorem 13 (Löwenheim–Skolem for graphs)**.** *Let $G = (V, E)$ be a graph with an infinite vertex set V. Let $\lambda \geq |\mathbb{N}|$ be any infinite cardinality. There is a graph G' with λ-many vertices such that $G \simeq G'$.*

This is a special case of a deep theorem of mathematical logic called the **Löwenheim–Skolem theorem**. Proving this theorem, even for graphs, is quite beyond the scope of this chapter. While the *upward* part of the theorem, that shows we can find a bigger infinite graph elementarily equivalent to G, the *downward* part is very tricky and requires a deep understanding of first order logic. However, we can try to give you a little intuition for why it's true. We already saw that the first order lens has a hard time separating finite graphs from infinite graphs. Recall that F_n was the first order property of having exactly n vertices. We asserted that "infinite-ness" is not a first order property. However, we can say that a graph G is infinite if and only if $G \vDash \neg F_n$ for all $n \in \mathbb{N}$. There might not be a single first order sentence corresponding to infinite-ness, but we can hold infinity in our hands with the infinite collection of sentences $\{\neg F_0, \neg F_1, \neg F_2, \cdots\}$. When it comes to distinguishing different infinite cardinalities on the other hand, we can't even grope at the concept. Since every first order sentence is finite in length and can quantify over at most finitely many vertices of a graph, there isn't even an infinite family of sentences that can distinguish countable from uncountable graphs. There is simply no way to tell infinite cardinalities apart.

Since the first order world lacks the tools to distinguish infinite cardinalities, given an infinite graph G, there is no fundamental barrier to having elementarily equivalent graph G' of a different infinite size. And indeed, by carefully picking our vertices and edges, we can always show that such graphs exists. This theorem has an important corollary.

Corollary 14. *Let G and G' be graphs. If G and G' both have the extension property, then G and G' are elementarily equivalent.*

PROOF: It suffices to show that if G is a graph with the extension property, then G is elementarily equivalent to the Rado graph R. Note that G has the extension property if and only if $G \vDash E_{n,m}$ for all $n, m \in \mathbb{N}$. By the Löwenheim–Skolem theorem for graphs, there is a countable graph C such that $G \simeq C$. Thus C satisfies the (n, m)-extension property for all $n, m \in \mathbb{N}$, so C has the extension property. By Theorem 9, C must be isomorphic to the Rado graph R. Isomorphism implies elementary equivalence, ergo $G \simeq C \simeq R$. □

Maybe now you're starting to get an inkling of where we're going. Suppose G is a graph with the extension property. Even though G could have any infinite cardinality, and there are structural properties (like having a clique of certain infinite size) that G may-or-may-not have, every *first order property* of G is fully determined. Once you know that G has the extension property, you know every single first order sentence ϕ such that $G \vDash \phi$. Namely, they're exactly the same first order sentences that the Rado graph R satisfies.

Completeness and entailment

We take a moment to recast Corollary 14 in slightly different language. We aren't going to prove anything new in this subsection, but merely look at this result from a different perspective.

Let I be a (possibly infinite) index set, and let $T := \{\phi_i : i \in I\}$ be a collection of first order sentences. In the world of mathematical logic, T is commonly called a **first order theory**, but we will try to avoid this language so as to avoid overburdening you with new nouns. Let Ψ be some (not necessarily first order) property. We say that Ψ is **entailed** by T if whenever a graph $G \vDash \phi_i$ for all $i \in I$, we also have that Ψ is true of G. Let's look at a few examples before we proceed.

Example 15. *Let $T := \{\phi_i : i \in I\}$. Each ϕ_i is entailed by T.*

Example 16. *Let ϕ be a first order property, and let $T := \{\phi, \neg\phi\}$. Every property Ψ is entailed by T.*

There is no graph G such that $G \vDash \phi$ and $G \vDash \neg\phi$. Consequently, every property Ψ is entailed vacuously. This is commonly known as the **semantic principle of explosion**.

Example 17. *Recall that F_n denotes the first order sentence that denotes having exactly n vertices. Let $T := \{\neg F_0, \neg F_1, \neg F_2, \cdots\}$. The property of having infinitely many vertices is entailed by T.*

Example 18. *Connectedness is entailed by $E_{2,0}$.*

Sometimes, a collection of sentences $T := \{\phi_i : i \in I\}$ will entail every first order property or its negation. That is, for every first order property ψ, either ψ or $\neg\psi$ is entailed by T. In this case, we call T **complete**.

Note that for T to be complete does *not* require that $\psi \in T$ or $\neg\psi \in T$ for every first order ψ. For example, we've already seen that $\{\phi, \neg\phi\}$ is complete. Let's look at a more interesting example.

Theorem 19 (T_{Ex} complete)**.** *$T_{Ex} := \{E_{n,m} \ : \ n, m \in \mathbb{N}\}$ is complete.*

This is just a reframing of Corollary 14. If a graph G satisfies $E_{n,m}$ for all $n, m \in \mathbb{N}$, then G has the extension property. It follows that G is elementarily equivalent to the Rado graph, so all first order properties are entailed by T_{Ex}.

Compactness

We end this section with one last important aspect of first order properties. A few pages ago, we gave a bunch of examples of structural properties of graphs that are not first order. You might have noticed that we didn't actually *prove* that these properties aren't first order. We had two reasons for doing so. First, we believe it's far more important for you, the reader, to walk away from this chapter with an internal sense of what kinds of properties aren't first order than it is for you to walk away knowing how to prove something isn't first order. Second, while it's easy to prove that a property is first order (just write down a sentence!), it's much harder to prove a property isn't first order. How do you know there isn't some clever way to write down "G has infinitely many vertices" as a first order sentence? Clearly you can't — the first order language isn't expressive enough. But there are infinitely many first order sentences. How would you prove there isn't some sentence, with hundreds of trillions of quantifiers, that somehow manages to be satisfied by G if and only if G has infinitely many vertices? There's a standard trick.

Theorem 20 (Compactness for graphs)**.** *Let ψ be a first order property entailed by a collection of first order sentences T. There is a finite subset $T_0 \subseteq T$ such that ψ is also entailed by T_0.*

Like Theorem 13, this is another special case of a deep result in mathematical logic called the **compactness theorem**. Before we discuss why this theorem is true, let's look at some consequences. As we promised, compactness can be used to prove that certain properties are not first order properties. For example, let's look at infinite-ness and the extension property.

Proposition 21. *There is no first order sentence ϕ_∞ such that $G \vDash \phi_\infty$ if and only if G is an infinite graph.*

Proof: For sake of contradiction, suppose there was such a ϕ_∞. Recall that F_n denotes the first-order sentence that says "G has exactly n vertices." The sentence ϕ_∞ would be entailed by the collection of sentences $\{\neg F_0, \neg F_1, \neg F_2, \cdots\}$. By compactness, there would be some finite subcollection $\{\neg F_{n_1}, \cdots, \neg F_{n_k}\}$ that also entails ϕ_∞. But these sentences are satisfies by finite graphs that don't have n_j-many vertices for each $j = 1, \cdots, k$. This is a contradiction, so ϕ_∞ cannot exist. □

Proposition 22. *There is no first order sentence ϕ_{Ex} such that $G \vDash \phi_{Ex}$ if and only if G has the extension property.*

Proof: This is barely harder than the previous example. If there were such a first order sentence ϕ_{Ex}, it would be entailed by T_{Ex}, and by some finite subset $\{E_{n_1,m_1}, \cdots, E_{n_k,m_k}\} \subseteq T_{\text{Ex}}$. But it's easy to see that there are finite graphs — which necessarily cannot have the extension property — that satisfy E_{n_j,m_j} for each $j = 1, \cdots, k$. The the extension property is not a first order property. □

Similar arguments can be made to show that connectedness and being cycle-free (AKA being a tree) are not first order properties.

Okay, so why is the compactness theorem true? Unlike the Löwenheim-Skolem theorem, we will actually sketch a proof of the compactness theorem for graphs. But first, some history. Originally, the compactness theorem for first order logic was viewed as a corollary to **Gödel's completeness theorem**. That's not a typo — in addition to the famous incompleteness theorem, Gödel also proved a result called the completeness theorem. Yes it's confusing nomenclature. No these theorems don't contradict each other. Anyway, the classical argument for the compactness theorem goes as follows. If ψ is a first order property entailed by a collection of first order sentences T, then ψ can be *proved* from T. This is no different than how in a first linear algebra class, you proved that all vector spaces (read: things that satisfy the collection of sentences corresponding to the vector space axioms) have a certain property by writing a proof that that property can be proven from the vector space axioms. The fact that you can always find a proof of the entailment is the crux of Gödel's completeness theorem. Now that proof of ψ is, in essence, a finite list of logical deductions. Each deduction can appeal to at most a finite number of sentences in T, so in total, only finitely many properties from T are used in the proof. Gather up these finitely many sentences, and put them in a box called T_0. Bam. ψ is also entailed by the finite set T_0. Compactness is proven.

This proof isn't very satisfying. First, unless we now separately justify the completeness theorem (really? *every* entailment can be proven? And what exactly do we mean by a "finite list of deductions"?), all we've done is replace one foundational theorem of mathematical logic with another. Second, the compactness theorem is, at its core, a semantic theorem. The compactness theorem for graphs is about what is true about graphs, not about what we can prove about graphs from the definition of a graph. Compactness simply has nothing to do with proofs, syntax, and deduction *per se*. This is the perspective taken by contemporary model theorists, and they have devised a multitude of purely semantic proofs of the compactness theorem. We sketch one such proof in Section 7.

6 The Zero-One Law for Finite Graphs

We finally arrive at the pièce de résistance of the chapter — the zero-one law for finite graphs. Here's the idea. Let $G(N)$ denote a random graph on N vertices, with each edge generated independently with a fair coinflip[1]. Earlier in this chapter, when we analyzed the infinite random graph game we saw that

$$\mathbb{P}[G(\mathbb{N}) = R] = 1$$

where we understand $G(\mathbb{N})$ to be the random graph on countably many vertices. For all finite N, though, a panopoly of different random graphs are possible. How likely is $G(N)$ to have a certain property, and how does this change as N changes? Let's look at a few examples. We already have two:

Example 23. *Let μ_N and ν_N denote the probabilities that $G(N)$ has an even number of edges, and at least $\frac{1}{2}\binom{N}{2}$ edges respectively. As $N \to \infty$, both μ_N and ν_N converge to $\frac{1}{2}$.*

Both of these properties are not first order. This is trivial for the property of having at least $\frac{1}{2}\binom{N}{2}$ edges (this requires an appeal to the number of vertices), but the other is a little trickier. Hint: it follows yet again from compactness. So what if we focus in on first order properties? A few more examples will help.

[1]everything in this section, including the zero-one law can be extended random graphs generated by coins with arbitrary probability $0 < p < 1$, but we stick to fair coins for simplicity.

Example 24. *Let μ_N denote the probability that $G(N)$ is a complete graph. As $N \to \infty$, we have $\mu_n \to 0$.*

The probability that $G(N)$ is complete is given exactly by $\mu_N = (1/2)^{\binom{N}{2}}$. As $N \to \infty$, this quantity goes to zero. We capture this statement about density by saying that $G(N)$ is **asymptotically almost never** complete.

Example 25. *Let μ_N denote the probability that $G(N)$ has the (r,s)-extension property for fixed $r, s \in \mathbb{N}$. As $N \to \infty$, we have $\mu_N \to 1$.*

Suppose $N > a + b$. For disjoint collections of vertices A, B with $|A| = r$ and $|B| = s$, let v be a vertex of $G(N)$ not in $A \cup B$. The probability that v is a witness in $\mathrm{Ex}(A, B)$ is exactly $(1/2)^{r+s}$. It follows that the probability that $\mathrm{Ex}(A, B)$ is empty is given by

$$\mathbb{P}[\mathrm{Ex}(A,B) = \varnothing] = \left(1 - (1/2)^{r+s}\right)^{N-(r+s)}.$$

There are $\binom{N}{r}\binom{N-r}{s}$ ways to pick such vertex sets A and B, so we see

$$\mathbb{P}[G(N) \not\models E_{r,s}] \leq \binom{n}{r}\binom{n-r}{s}\left(1 - (1/2)^{r+s}\right)^{N-(r+s)}.$$

As $N \to \infty$, the expression on the right hand side goes to zero. It follows that as $N \to \infty$, the probability that $G(N)$ satisfies $E_{r,s}$ goes to 1. We say that $G(N)$ **asymptotically almost surely** has the (r,s)-extension property.

Both of these examples of first order properties have "trivial" densities of zero and one. If you play around and compute the densities of other first order properties ϕ, you will surely find that like in the examples above, you always get convergence to zero or one. The property of having a specific induced subgraph? Asymptotically almost surely happens. The property of having an isolated vertex or a universal vertex? Asymptotically almost never. It always plays out in one of these two ways. This observation leads us to the **zero-one law for finite graphs**.

Theorem 26 (Zero-one law)**.** *Let ϕ be a first order property. $G(N)$ satisfies ϕ asymptotically almost surely or asymptotically almost never.*

There's also a simple procedure for determining whether ϕ holds asymptotically almost surely or almost never: simply check if the property is true of the Rado graph R. If $R \models \phi$, then $G(N)$ asymptotically almost surely satisfies ϕ. Don't get confused here — this procedure only works for first order properties ϕ. There are non-first order properties of R that do not hold asymptotically almost surely for $G(N)$. To pick an easy example, the property of having an infinite clique.

To head off another potential point of confusion, the converse to the zero-one law isn't true. That is, if ϕ is a property that holds asymptotically almost surely or almost never, it does not follow that ϕ is first order. For example, $G(N)$ is asymptotically almost surely connected. If a graph satisfies the $(2,0)$-extension property, it has to be connected. Since the $(2,0)$-extension property holds asymptotically almost surely, it follows that $\lim_{N\to\infty}$, we have $\mathbb{P}[G(N) \text{ is connected}] \to 1$. Nonetheless, compactness can be used to show that connectivity isn't a first order property.

Proof of theorem 26: Believe it or not, we've already done almost all of the work to prove the zero-one law for graphs. We've already given you all the pieces — they just need to be assembled. To start, let ϕ be a first order property. We have two cases to consider: whether $R \models \phi$, or $R \not\models \phi$.

We first consider the case where $R \vDash \phi$. We must show that ϕ holds asymptotically almost surely for $G(N)$. Completeness (Corollary 19) tells us that ϕ is entailed by the collection of first order properties $T_{\text{Ex}} := \{E_{n,m} : n, m \in \mathbb{N}\}$. By compactness (Theorem 20), ϕ is entailed by a finite subcollection $\{E_{n_1,m_1}, \cdots, E_{n_k,m_k}\} \subseteq T_{\text{Ex}}$. And from Example 25, we know for each $j = 1, \cdots, k$, that E_{n_j,m_j} holds asymptotically almost surely for $G(N)$. Some basic probability proves that the conjunction

$$\psi := E_{n_1,m_1} \wedge E_{n_2,m_2} \wedge \cdots \wedge E_{n_k,m_k}$$

holds asymptotically almost surely. But for any graph G, if $G \vDash \psi$, then $G \vDash \phi$. So $1 \geq \mathbb{P}[G(N) \vDash \phi] \geq \mathbb{P}[G(N) \vDash \psi]$. We already proved that ψ holds asymptotically almost surely, so ϕ holds asymptotically almost surely as well.

We now consider the case that $R \nvDash \phi$. This is covered by the first case with just a small sleight of hand. Simply observe that $R \nvDash \phi$ if and only if $R \vDash \neg\phi$. Applying the first case to $\neg\phi$ tells us that that $G(N) \vDash \neg\phi$ asymptotically almost surely. Consequently, $G(N) \vDash \phi$ asymptotically almost never. This proves the zero-one law for finite graphs. $\square$

The zero-one law is a bittersweet theorem. On one hand, the argument is stunningly beautiful. We're proving a statement about finite graphs. To do so, we pass from the finite into the world of infinite graphs. We show that the Rado graph is, in a certain sense, the limit of randomly chosen finite graphs. And then we use the tools of mathematical logic to send truths about the Rado graph back to truths about finite graphs! The proof boomerangs through infinity and back into the finite. The zero one law is a seemingly magical property of finite graphs.

But there's also a pessimistic view. The zero-one law proves that no first order property could be a good bet in the finite random graph game. Earlier, we remarked that most natural, interesting structural properties of graphs are not first order properties. The zero-one law should make this feeling visceral. The vocabulary of the first order world is so impoverished that it can't even make a statement that isn't true of almost every finite graph or almost no finite graph. It is a scathing indictment of the first order language.

7 More on compactness

We still owe you a proof of the compactness theorem. Compactness is more than the crux of the proof of the zero-one law for graphs. It is arguably the fundamental theorem of model theory, and myriad theorems follow from it. The majority of working mathematicians do not know this theorem, and that is simply a travesty. However, we accept that the material in this section is a little tricker than what has come before, and might not be to everyone's tastes.

Before we prove the compactness theorem, we need two things: an alternative characterization of the compactness theorem, and a certain construction that takes in some small graphs, and spits out a big one. Let's start with the alternative characterization.

Let T be a collection of first order sentences about graphs. We say that T is **satisfiable** if there is a graph G such that $G \vDash \phi$ for all $\phi \in T$. Alternatively, we say that T is **finitely satisfiable** if for every finite subset $T_0 \subseteq T$, there is a graph G that satisfies T_0. Clearly if T is satisfiable, then it is finitely satisfiable. But what about the converse? Well...

Lemma 27. *Let T be a collection of first order sentences about graphs. The following are equivalent:*

1. *If ϕ is entailed by T, then there is a finite $T_0 \subseteq T$ such that ϕ is entailed by T_0;*
2. *If T is finitely satisfiable, then T is satisfiable.*

The first statement, of course, is the compactness theorem for graphs (Theorem 20). Let's prove the lemma.

PROOF: Suppose that the compactness theorem holds, and that T is not satisfiable. We will prove that T is not finitely satisfiable. Let $\perp$ denote any contradictory first order sentence. For example, we could take $\perp := \exists x(x \neq x)$. Since T is not satisfiable, it vacuously entails $\perp$. By the compactness theorem, there is a finite $T_0 \subseteq T$ such that $T_0 \vDash \perp$. But since no graph G can satisfy the contradictory $\perp$, it follows that T_0 is not satisfiable, and that T is not finitely satisfiable. This proves the first direction.

Conversely, suppose T is satisfiable whenever T is finitely satisfiable, and that ϕ is entailed by T. The collection $T \cup \{\neg\phi\}$ cannot be satisfiable. By our assumption, this means that there must also be some finite collection $T_0 \cup \{\neg\phi\} \subseteq T \cup \{\neg\phi\}$ which is also not satisfiable. It immediately follows (possibly by the semantic principle of explosion if T_0 itself is not satisfiable) that ϕ must be entailed by T_0. This proves the equivalence. □

This alternative version of the compactness theorem further demonstrates the semantic nature of compactness. While less useful for proving the zero-one law for graphs, this version of compactness is, in some sense, what's really going on under the hood.

Now it's time for a weird construction on graphs. Some parts of this are going to feel somewhat miraculous unless you're well acquainted with ultrafilters. Let I be a (possibly infinite) set, and let $\mathcal{U}$ be a collection of subsets of I with the following properties:

1. (non-trivial) $\varnothing \notin \mathcal{U}$;
2. (monotone) If $A \in \mathcal{U}$ and $B \supseteq A$, then $B \in \mathcal{U}$;
3. (finite intersection) If $A, B \in \mathcal{U}$, then the intersection $A \cap B \in \mathcal{U}$;
4. (maximal) For every $A \subseteq I$, either $A \in \mathcal{U}$ or its complement $(I - A) \in \mathcal{U}$.

Now for each $i \in I$, fix a graph G_i. We fix a new graph, $\prod_{\mathcal{U}} G_i$, according to the following specifications. A vertex $v \in \prod_{\mathcal{U}} G_i$ is a choice of a vertex $v_i \in G_i$ for each $i \in I$. That is, $v = \{v_i \in G_i\}_{i \in I}$. We say two vertices $v = \{v_i\}_{i \in I}$ and $w = \{w_i\}_{i \in I}$ are equal if and only if $\{i \in I \,:\, v_i = w_i\} \in \mathcal{U}$. Formally, we need to mod out by this relation, but let's not get overburdened with notation so late in this chapter. Similarly, we have an edge $v \sim w$ in $\prod_{\mathcal{U}} G_i$ if and only if $\{i \in I \,:\, v_i \sim w_i\} \in \mathcal{U}$.

Since $\sim$ is symmetric and antireflexive on each underlying graph G_i, we get a valid graph structure. This graph is called the **ultraproduct** of the graphs $\{G_i\}_{i \in I}$ with respect to the **ultrafilter** $\mathcal{U}$. We care about this construction because it has the following important property:

Proposition 28. *Let ϕ be a first order property of a graph. If $\{i \in I \,:\, G_i \vDash \phi\} \in \mathcal{U}$, then $\prod_{\mathcal{U}} G_i \vDash \phi$.*

This can be proven by simply following your nose. The first order property ϕ is ultimately made up from appeals to the equality and adjacency relations, plus logical connectives and quantifiers. We specifically chose for $\prod_{\mathcal{U}} G_i \vDash \phi$ to have equalities and edges based on having the equality/edge on a collection of indices in $\mathcal{U}$. Since $\{i \in I \,:\, G_i \vDash \phi\} \in \mathcal{U}$, the requisite equalities and adjacencies are guaranteed to hold.

Finally, we're equipped to prove the compactness theorem for graphs.

Proof of the compactness theorem for graphs: Let T be a finitely satisfiable collection of first order properties of graphs. By Lemma 27, to prove compactness, it will suffice to show that T is satisfiable. To that end, let $\mathcal{F}$ denote the collection of all non-empty finite subsets of properties from T. For each $T_0 \in \mathcal{F}$, let T_0^* denote the collection of all finite collections of properties in T that extend T_0. That is,

$$T_0^* := \{S_0 \in \mathcal{F} \,:\, T_0 \subseteq S_0\}.$$

Consider the set $\mathcal{U}_0 := \{T_0^* \,:\, T_0 \in \mathcal{F}\}$. Is this an ultrafilter on $\mathcal{F}$? Not quite, but it can be extended to one using some set theoretic nonsense. It's quite easy to see that $\mathcal{U}_0$ is closed under intersections, since $T_0^* \cap R_0^* = (T_0 \cup R_0)^*$ for all $T_0, R_0 \in \mathcal{F}$. That's all it takes to ensure that there is an ultrafilter $\mathcal{U}$ extending $\mathcal{U}_0$.

We assumed that T is finitely satisfiable. For each $T_0 \in \mathcal{F}$, there is a graph $G_{T_0} \vDash F$. Now we take the ultraproduct $G := \prod_{\mathcal{U}} G_F$. We claim that G satisfies the theory T. Indeed, for any $\phi \in T$, consider the set $\{\phi\}^* \in \mathcal{U}$. For every $F \in \{\phi\}^*$, we have that $\phi \in F$, and thus $G_F \vDash \phi$. By Proposition 28, $G \vDash \phi$. $\square$

Now that is a purely semantic proof of compactness.