6. The Main Theorem

We fix an irrational $\alpha \in (0,1)$ throughout this chapter. All probabilities are with respect to the random graph $G(n,p)$ with $p = n^{-\alpha}$. We recall the statement of our goal, the Main Theorem 1.4.1: For any first order A

$$\lim_{n\to\infty} \Pr\big[G(n, n^{-\alpha}) \models A\big] = 0 \text{ or } 1$$

6.1 The Look-Ahead Strategy

Our approach is through the Ehrenfeucht Game as described in Section 2.1. We fix the number k of moves. We shall give a strategy for Duplicator so that, as $n, m \to \infty$, she almost surely wins $\text{EHR}(G_1, G_2, k)$ where $G_1 \sim G(n, n^{-\alpha})$ and $G_2 \sim G(m, m^{-\alpha})$ are independently chosen.

Let $0 = t_0, t_1, \ldots t_{k-1}$ be nonnegative integers. The $(t_0, \ldots, t_{k-1})-$ *look-ahead strategy* for Duplicator is easy to describe. Duplicator makes any moves in response to Spoiler so that when there are i rounds remaining in the game the t_i-types of the vertices chosen are the same in both graphs. That is, if $x_1, \ldots, x_{k-i} \in G_1$, $y_1, \ldots, y_{k-i} \in G_2$ have been chosen then there is a graph isomorphism from $\text{cl}_{t_i}(x_1, \ldots, x_{k-i})$ to $\text{cl}_{t_i}(y_1, \ldots, y_{k-i})$ sending each x_j to its corresponding y_j.

Of course, it may well be that Duplicator is unable to keep to this strategy. In that case she loses. But if she is able to keep to this strategy then at the end of the game the 0-closures are the same and she has won. We shall give explicit (though surprisingly complicated) $t_0, \ldots, t_{k-1}$ so that Duplicator shall almost surely be able to keep to this strategy. Formally, we find t_i by induction on i. Note that as i represents the number of remaining moves we are really working backwards from the end of the game. We need show that almost surely for *every* $x_1, \ldots, x_{k-i} \in G_1$ and $y_1, \ldots, y_{k-i} \in G_2$ that have the same t_i-type and every $x_{k-i+1} \in G_1$ (Spoiler move) there exists $y_{k-i+1} \in G_2$ (Duplicator move) so that the resulting $k - i + 1$-tuples have the same t_{i-1}-type. [Of course, Spoiler could also move in G_2 but this case is the same by symmetry.] For convenience of exposition we consider the first ($i = k$) and final ($i = 1$) moves separately. In a formal sense this is unnecessary.

6.1.1 The Final Move

We set $t_1 = 1$. Assume $x_1, \ldots, x_{k-1} \in G_1$, $y_1, \ldots, y_{k-1} \in G_2$ have been chosen with the same 1-closure. Now Spoiler moves and by symmetry we can assume he picks $x_k \in G_1$. (Note that we cannot assume x_k is a random choice, quite the opposite!) Let w be the number of previously selected x's adjacent to the newly selected x_k.

Case 1 (Inside): $1 - w\alpha < 0$. Then $x_k \in \mathrm{cl}_1(x_1, \ldots, x_{k-1})$ since the rooted graph with $v = 1$ nonroot and w edges is rigid. Let $\Psi\colon \mathrm{cl}_1(x_1, \ldots, x_{k-1}) \to \mathrm{cl}_1(y_1, \ldots, y_{k-1})$ be the isomorphism guaranteed by the 1-types being the same. Spoiler selects $y_k = \Psi(x_k)$. (Wily Spoiler's attempt to trick Duplicator, as in Section 2.5, is thwarted by her having looked ahead and assured that not only the induced graphs but the 1-closures were identical.)

Case 2 (Outside): $1 - w\alpha > 0$. The rooted graph with $k - 1$ roots and one nonroot adjacent to w of the roots is now safe. By our Generic Extension Theorem 5.3.1 almost surely for every $k - 1$ vertices in G_2 there is a vertex adjacent to any prescribed w of them and no others. Duplicator picks that $y_k \in G_2$ adjacent to just those w of the $y_j \in G_2$ such that the x_k selected by Spoiler was adjacent in G_1 to x_j.

6.1.2 The Core Argument (Middle Moves)

Let us fix i with $1 \le i \le k - 1$ and let t_i be given. We select t_{i+1} so that

1. $t_{i+1} \ge t_i$
2. Almost surely the t_i-closure of any $k - i$ vertices has at most $t_{i+1} - 1$ nonroots.

The existence of such t_{i+1} is a consequence of the Finite Closure Theorem 4.3.2. For notational convenience we set $t = t_i$, $u = t_{i+1}$. Assume $x_1, \ldots, x_{k-i-1} \in G_1$, $y_1, \ldots, y_{k-i-1} \in G_2$ have been chosen with the same u-closure. Set $\mathbf{x} = (x_1, \ldots, x_{k-i-1})$, $\mathbf{y} = (y_1, \ldots, y_{k-i-1})$ for further notational convenience. Let $\Psi\colon \mathrm{cl}_u(\mathbf{x}) \to \mathrm{cl}_u(\mathbf{y})$ be the graph isomorphism showing that their u-types are the same. Now Spoiler selects $x \in G_1$.

Case 1 (Inside) $x \in \mathrm{cl}_u(\mathbf{x})$. Spoiler selects $y = \Psi(x)$. As $t \le u$ the u-closure of $\mathbf{x}, x$ is contained in $\mathrm{cl}_u(\mathbf{x})$ (and also for $\mathbf{y}$) so that restricting Ψ gives an isomorphism from $\mathrm{cl}_t(\mathbf{x}, x)$ to $\mathrm{cl}_t(\mathbf{y}, y)$.

Case 2 (Outside) $x \notin \mathrm{cl}_u(\mathbf{x})$. Set $H = \mathrm{cl}_t(\mathbf{x}, x)$ and $R = \mathrm{cl}_t(\mathbf{x}, x) \cap \mathrm{cl}_u(\mathbf{x})$ and consider the rooted graph (R, H). As $x \in H$ but by assumption $x \notin R$ this is a legitimate extension.

We claim (R, H) is safe. Otherwise by Property 4.1.7 it would have a rigid subextension (R, H'). The number of nonroots of (R, H') would be at most the number of nonroots of $(\mathbf{x}, \mathrm{cl}_t(\mathbf{x}, x))$ which is one plus the number of nonroots of $\big((\mathbf{x}, x), \mathrm{cl}_t(\mathbf{x}, x)\big)$. We've designed $u = t_{i+1}$ so that this is at most $1 + (u - 1) = u$. Then

$$H' \subseteq \mathrm{cl}_u(R) \subseteq \mathrm{cl}_u(\mathrm{cl}_u(\mathbf{x})) = \mathrm{cl}_u(\mathbf{x}) \subseteq R$$

a contradiction.

Set $R' = \Psi(R)$. We apply Generic Extension Theorem 5.3.1 to give a u-generic (R, H) extension over R', call it (R', H'). The isomorphism Ψ, limited to $R \to R'$, extends to a graph isomorphism $\Psi^+ \colon H \to H'$. Spoiler selects $y = \Psi^+(x)$.

Does this work? As H, H' are isomorphic $\mathrm{cl}_t(\mathbf{y}, y)$ certainly contains H'. There are no additional edges in H' since there were none in R' and the extension was generic. Can there be more points in $\mathrm{cl}_t(\mathbf{y}, y)$? Set $\mathrm{N\scriptstyle EW}' = \mathrm{cl}_t(\mathbf{y}, y) - H'$. If $\mathrm{N\scriptstyle EW}' \neq \emptyset$ then $\mathrm{cl}_t(\mathbf{y}, y)$ would be a rigid extension over H' and, by the general bound, $\mathrm{N\scriptstyle EW}'$ would have at most u vertices. Thus the $v \in \mathrm{N\scriptstyle EW}'$ would be adjacent only to vertices in $H' \cap \mathrm{cl}_u(\mathbf{x})$ (not to $H' - \mathrm{cl}_u(\mathbf{x})$) and would be a rigid extension over $H' \cap \mathrm{cl}_u(\mathbf{x})$. Hence $\mathrm{N\scriptstyle EW}' \subset \mathrm{cl}_u\big(\mathrm{cl}_u(\mathbf{x})\big)$ which is simply $\mathrm{cl}_u(\mathbf{x})$. But then back in G_1, setting $\mathrm{N\scriptstyle EW} = \Psi^{-1}(\mathrm{N\scriptstyle EW}')$ we have the same picture with $H \cup \mathrm{N\scriptstyle EW}$ isomorphic to $H' \cup \mathrm{N\scriptstyle EW}'$. In G_2 all points of $H' \cup \mathrm{N\scriptstyle EW}'$ can be reached from $\mathbf{y}, y$ by rigid extensions of at most t nonroots and those extensions never go outside of $H' \cup \mathrm{N\scriptstyle EW}'$. But then the same would be true in G_1 with the isomorphic $H \cup \mathrm{N\scriptstyle EW}$ and that would give $\mathrm{cl}_t(\mathbf{x}, x)$ extra vertices that it doesn't have.

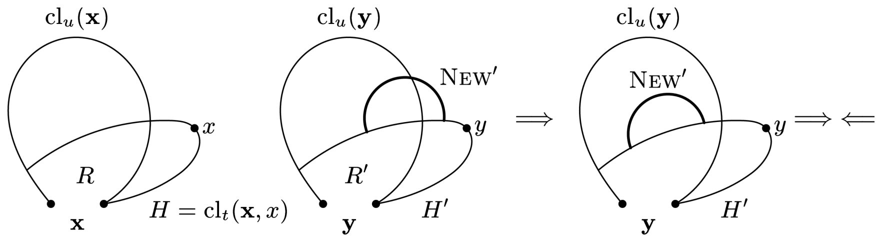

Fig. 6.1. Spoiler plays Outside x. Duplicator finds y with u-generic (R, H)-extension. If $\mathrm{cl}_t(\mathbf{y}, y)$ had other vertices $\mathrm{N\scriptstyle EW}'$ the genericity would force them inside $\mathrm{cl}_u(\mathbf{y})$ but then they would have been in $\mathrm{cl}_t(\mathbf{x}, x)$ as well

6.1.3 The First Move

Set $t = t_k$. On the first move Spoiler selects some $x \in G_1$. Duplicator calculates $\mathrm{cl}_t(x)$ and must find an $y \in G_2$ with the same t-type. To show that she almost surely succeeds one needs that every t-type of a single vertex either almost surely or almost never appears in $G(n, n^{-\alpha})$. There are only a finite number (by the Finite Closure Theorem 4.3.2) of possible t-types to consider so it suffices to show this for any particular one. We write the t-type as the graph $H = \mathrm{cl}_t(x)$, with vertex x specified. We look only at logically possible H, so that $(\{x\}, H)$ is a rigid extension.

Suppose H contains a subgraph H_1 with v_1 vertices, e_1 edges where $v_1 - e_1\alpha < 0$. The expected number of copies of H_1 is $O(n^{v_1}p^{e_1}) = O(n^{v_1 - e_1\alpha})$

which is $o(1)$ so almost surely there are no copies of H_1 and hence no copies of H and hence no x with $H = \text{cl}_t(x)$ exists.

Otherwise $(\emptyset, H)$ is a safe extension. Then not only does there exist a copy of H but by the Generic Extension Theorem 5.3.1 there exists an induced copy of H which is t-generic over the empty set – which means that $\text{cl}_t(H) = H$. We know $\text{cl}_t(x)$ contains H, but it is also contained in $\text{cl}_t(H) = H$ and therefore it is precisely H.

6.2 The Original Argument

We begin by restating the crucial idea of Section 6.1.2 which in some sense is the centerpiece of the entire argument.

Theorem 6.2.1. *Let $u \geq t$ be such that almost surely the t-closure of any $k+1$ vertices has at most $u-1$ nonroots. Let H be any possible value of* $\text{cl}_u(\mathbf{x})$, *where we set* $\mathbf{x} = (x_1, \ldots, x_k)$. *Let H_1 be any possible value of* $\text{cl}_t(\mathbf{x}, x)$. *Then almost surely either*
• *For every $\mathbf{x}$ with* $\text{cl}_u(\mathbf{x}) \cong H$ *there exists x with* $\text{cl}_t(\mathbf{x}, x) \cong H_1$
or
• *For every $\mathbf{x}$ with* $\text{cl}_u(\mathbf{x}) \cong H$ *there does* not *exist x with* $\text{cl}_t(\mathbf{x}, x) \cong H_1$.

We want H, H_1 to have common vertices $x_1, \ldots, x_k$. They may or may not have other common vertices. We call H^* a *picture* if it is derived from H, H_1 by identifying the roots (in the prescribed order) and identifying some (possibly none) other pairs of vertices and otherwise keeping the vertices distinct. As H^* has bounded size there are only a finite number of possible pictures H^*. We actually show that for every such H^* almost surely either
• For every $\mathbf{x}$ with $\text{cl}_u(\mathbf{x}) \cong H$ there exists x with $\text{cl}_t(\mathbf{x}, x) \cong H_1$ and $H \cup H_1 \cong H^*$
or
• For every $\mathbf{x}$ with $\text{cl}_u(\mathbf{x}) \cong H$ there does not exist x with $\text{cl}_t(\mathbf{x}, x) \cong H_1$ and $H \cup H_1 \cong H^*$.

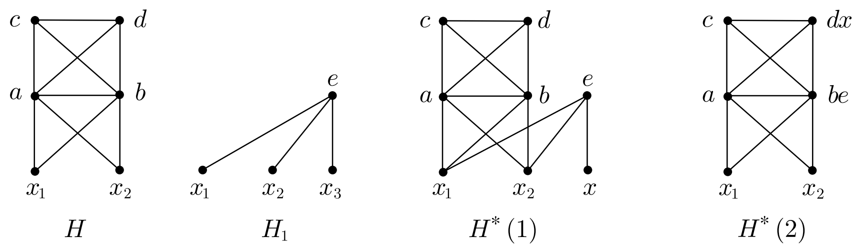

Fig. 6.2. $\alpha = \pi/7$. $H \cong \text{cl}_2(x_1, x_2)$. $H_1 \cong \text{cl}_1(x_1, x_2, x)$. $H^*(1)$, $H^*(2)$ are two (of many) possible pictures given by identification. In $H^*(2)$ x is Inside. But in $H^*(2)$ $a, c \in \text{cl}_1(x_1, x_2, x)$, hence no x can exist with this picture. In $H^*(1)$ x is Outside. $\{x, e\}$ is safe over $H_0 = \{x_1, x_2\}$. By Generic Extension such x, e almost surely exist

In the picture H^* set $H_0 = H \cap H_1$. There are two cases.

Case 1 (Inside) $x \in H_0$. Then $H = \mathrm{cl}_u(\mathbf{x})$ determines $\mathrm{cl}_t(\mathbf{x}, x)$ so we must have $H_1 \subseteq H$ and can check if H_1 is indeed $\mathrm{cl}_t(\mathbf{x}, x)$.

Case 2 (Outside) $x \notin H_0$. Then we must have (H_0, H_1) safe. Otherwise by Property 4.1.7 it would have a rigid subextension (H_0, H_2) whose number of nonroots would be at most $u - 1 + 1 = u$ and then H_2 would have to be contained in H. It further must be that in H^* the t-closure of $\mathbf{x}, x$ is H_1, nothing more or less. But suppose these are satisfied. By the Generic Extension Theorem 5.3.1 for all $\mathbf{x}$ with $\mathrm{cl}_u(\mathbf{x}) \cong H$ there will exist a u-generic extension giving H^*. If the t-closure of $\mathbf{x}, x$ was H_2, strictly more than H_1, then H_2 would be a rigid extension over H_1 but by genericity it would be a rigid extension over H_0 but then it would be in H which we have already checked. □

The original proof of the Main Theorem 1.4.1 did not use the Ehrenfeucht game. Rather, it was an induction on the length of the statement. To make the induction go, however, we need to prove a statement for all predicates $P(x_1, \ldots, x_k)$ with any number k of free variables. Sentences have $k = 0$ free variables. As before, α is a fixed irrational number between zero and one.

Theorem 6.2.2. *For every predicate $P(x_1, \ldots, x_k)$ there exists a nonnegative integer t so that the following holds almost surely in $G(n, p)$ with $p = n^{-\alpha}$: for each t-type H either all $x_1, \ldots, x_k$ with that t-type satisfy P or no $x_1, \ldots, x_k$ with that t-type satisfy P.*

The proof is by induction on the length of the predicate P. For the atomic predicates $x_i \sim x_j$ and $x_i = x_j$ we can take $t = 0$ as the 0-closure includes this information. If the statement holds for P then it certainly holds for $\neg P$ with the same value of t. If the statement holds for P, Q with values t_1, t_2 then it holds for $P \wedge Q$ (or any Boolean function of P, Q) with the value $t = \max(t_1, t_2)$, as a t-type includes the information about the s-type for all smaller s. This leaves us with the one important case, a predicate of the form $Q = \exists_x P(x_1, \ldots, x_k, x)$. By induction there is a t so the $P(x_1, \ldots, x_k, x)$ holds if and only if the t-type of $x_1, \ldots, x_k, x$ is one of some finite list $H_1, \ldots, H_s$. Let u satisfy the conditions of Theorem 6.2.1. There is a finite list of potention u-types for $x_1, \ldots, x_k$. For each such u-type, either almost surely there exists x with $x_1, \ldots, x_k, x$ having t-type one of $H_1, \ldots, H_s$ or almost surely no such x exists. Call those u-types positive and negative respectively. Then almost surely Q holds if and only if the u-type of $x_1, \ldots, x_k$ is positive.