SHARP FUSS–CATALAN THRESHOLDS IN GRAPH BOOTSTRAP PERCOLATION

ZSOLT BARTHA, BRETT KOLESNIK, GAL KRONENBERG, AND YUVAL PELED

ABSTRACT. We study graph bootstrap percolation on the Erdős–Rényi random graph $\mathscr{G}_{n,p}$. For all $r \geqslant 5$, we locate the sharp K_r-percolation threshold $p_c \sim (\gamma n)^{-1/\lambda}$, solving a problem of Balogh, Bollobás and Morris. The case $r = 3$ is the classical graph connectivity threshold, and the threshold for $r = 4$ was found using strong connections with the well-studied 2-neighbor dynamics from statistical physics. When $r \geqslant 5$, such connections break down, and the process exhibits much richer behavior. The constants $\lambda = \lambda(r)$ and $\gamma = \gamma(r)$ in p_c are determined by a class of $\left(\binom{r}{2} - 1\right)$-ary tree-like graphs, which we call K_r-tree witness graphs. These graphs are associated with the most efficient ways of adding a new edge in the K_r-dynamics, and they can be counted using the Fuss–Catalan numbers. Also, in the subcritical setting, we determine the asymptotic number of edges added to $\mathscr{G}_{n,p}$, showing that the edge density increases only by a constant factor, whose value we identify.

1. INTRODUCTION

Bollobás [17] introduced the concept of *weak saturation* in graph theory, which gives rise to a process called *H-bootstrap percolation* (or *H-dynamics*) on a graph G. Given a fixed *template* graph H and an initial graph $G \subseteq K_n$, at each step of the process a new edge e is added, provided there is a copy of H in K_n for which e is the only edge that has not yet been added to G. If no such edge exists, the process terminates, and we denote the final graph by $\langle G \rangle_H$. We say that G is *weakly H-saturated* (or that it *H-percolates*) if every edge is eventually added, that is, $\langle G \rangle_H = K_n$.

In this work, we find, for every $r \geqslant 5$, the sharp threshold probability $p_c(n, K_r)$ at which the Erdős–Rényi random graph $\mathscr{G}_{n,p}$ is likely to K_r-percolate. This solves, in a strong form, an open problem of Balogh, Bollobás and Morris [9], who located the threshold up to polylogarithmic factors and asked for the threshold up to constant factors. This result concerns one of the central problems in graph bootstrap percolation. Since the seminal

2010 *Mathematics Subject Classification.* 05C05; 05C35; 05C65; 05C80; 60K35; 68Q80.

Key words and phrases. bootstrap percolation; cellular automaton; critical threshold; Fuss–Catalan numbers; phase transition; random graph; slow aging; weak saturation.

work of Bollobás [17], the case $H = K_r$ has attracted particular attention, and following Balogh, Bollobás, and Morris [9], the precise behavior of the K_r-dynamics on $G = \mathscr{G}_{n,p}$ has remained open.

The classical model of *bootstrap percolation* was introduced in the statistical physics literature by Chalupa, Leath and Reich [20], and is one of the most well studied of all *cellular automata*. Such processes, as pioneered by Ulam [51] and von Neumann [53], are notable in that, although they evolve only by local rules, they can exhibit complicated global behavior of interest; see, e.g., Gardner's [28] introduction to Conway's *Game of Life*. In vertex bootstrap percolation on a graph G, each vertex is initially *infected* with probability p, and further vertices are infected by some local rule. For instance, and most commonly, by the *r-neighbor rule*, whereby a vertex with at least r infected neighbors becomes infected. In this context, *percolation* is the event that all vertices are eventually infected. There are many variations; see, e.g., the survey by Morris [43] and references therein.

Graph bootstrap percolation can be seen as a generalization of the classical r-neighbor model, where the template graph H encodes a more complicated update rule. Notably, the "infection" in graph bootstrap percolation spreads amongst the *edges* of some graph, rather than its vertices. The initialization at $\mathscr{G}_{n,p}$ is therefore quite natural, since in bootstrap percolation each site is assumed to be initially infected independently with probability p.

The study of $\mathscr{G}_{n,p}$ focuses on graph properties that occur *with high probability*, i.e., with probability tending to 1 as $n \to \infty$. A main theme in random graph theory, initiated by the seminal work of Erdős and Rényi [22], is the systematic study of threshold probabilities for monotone graph properties. A property P is said to have a *sharp threshold* at p_c if for every fixed $\varepsilon > 0$, with high probability, $\mathscr{G}_{n,(1+\varepsilon)p_c}$ has property P but $\mathscr{G}_{n,(1-\varepsilon)p_c}$ does not. For example, a classical result from [22] shows that graph connectivity has a sharp threshold at $p_c \sim (\log n)/n$. Observe that $\langle G\rangle_{K_3} = K_n$ if and only if G is connected, and therefore thresholds for H-percolation can be viewed as higher forms of the celebrated graph connectivity threshold.

1.1. **Main results.**

As in [9], we let

$$p_c(n,H) = \inf\{p > 0 : \mathbb{P}(\langle \mathscr{G}_{n,p}\rangle_H = K_n) \geqslant 1/2\}$$

denote the *critical H-percolation threshold*. Roughly speaking, at this point $\mathscr{G}_{n,p}$ is likely to H-percolate. One of the main results in [9, Theorem 1] shows that, for every $r \geqslant 3$,

$$p_c(n,K_r) = n^{-1/\lambda + o(1)}, \tag{1.1}$$

as $n \to \infty$, where

$$\lambda = \frac{\binom{r}{2} - 2}{r-2}\,. \tag{1.2}$$

Intuitively, $-1/\lambda$ is the critical exponent since once $p \gg n^{-1/\lambda}$, the number of edges added in the first round of the K_r-dynamics is much larger than the initial number of edges in $\mathscr{G}_{n,p}$.

The classical result of Erdős and Rényi [22] implies that $p_c(n,K_3) \sim (\log n)/n$, where, as usual, we write $a_n \sim b_n$ if $a_n/b_n \to 1$, as $n \to \infty$. When $r = 4$, the threshold was identified up to constants in [9], and the existence of a sharp threshold at $p_c(n,K_4) \sim (3n\log n)^{-1/2}$ follows by [3,4,38].

Our first main result establishes a sharp threshold for K_r-percolation, in all of the remaining cases $r \geqslant 5$. Furthermore, we locate the precise first-order asymptotics of $p_c(n,K_r)$. Let $\alpha_d = (d+1)^{d+1}/d^d$ denote the asymptotic growth rate of the dth Fuss–Catalan numbers (see Section 3.2), that is, the number of $(d+1)$-ary plane trees.

Theorem 1.1. *Fix $r \geqslant 5$. Let $\gamma > 0$ be the unique positive root of*

$$(r-2)!\gamma^{r-2} = \alpha_{\binom{r}{2}-2}. \tag{1.3}$$

Then, for every $\varepsilon > 0$,

$$\mathbb{P}(\langle \mathscr{G}_{n,p} \rangle_{K_r} = K_n) \to \begin{cases} 0, & \text{if } \gamma n p^{\lambda} = 1-\varepsilon; \\ 1, & \text{if } \gamma n p^{\lambda} = 1+\varepsilon, \end{cases}$$

as $n \to \infty$. In particular,

$$p_c(n,K_r) \sim (\gamma n)^{-1/\lambda}.$$

Problem 3 in [9] asks for $p_c(n,K_r)$ up to constant multiplicative factors. Theorem 1.1 gives, moreover, the precise leading-order asymptotics. In addition, Problem 2 in [9] asks for which H is the critical H-percolation threshold sharp, and, by Theorem 1.1, the threshold is sharp when H is a clique.

The distinction between the behavior of K_r-percolation for $r = 4$ and for $r \geqslant 5$ is already evident from the absence of a polylogarithmic factor in $p_c(n,K_r)$. There are several interesting reasons for this transition. First, although all cliques K_r are, in a certain sense, *balanced*, only once $r \geqslant 5$ are they *strictly balanced* (see Definition 3.17 below). Related to this, on $\mathscr{G}_{n,p}$ the K_4-dynamics are essentially dominated by the simpler 2-neighbor bootstrap percolation dynamics. Specifically, the most likely way for $\mathscr{G}_{n,p}$ to K_4-percolate is to start with the two vertices in some edge $e \in E(\mathscr{G}_{n,p})$ and then iteratively add vertices with at least two neighbors amongst those previously added. By induction, after each such vertex is added, a K_4-percolating graph is obtained. In this sense, K_4-percolation on $\mathscr{G}_{n,p}$ (like most forms of bootstrap percolation) spreads by "nucleation," meaning that the percolation spreads locally from the boundary of a growing structure (the "nucleus") until it reaches a critical size (sometimes called a "critical droplet"), at which point the process transitions to a final stage of explosive growth. On the other hand, for $r \geqslant 5$, the connection to the $(r-2)$-neighbor dynamics breaks down, and K_r-percolation on $\mathscr{G}_{n,p}$ exhibits qualitatively different behavior. In particular, when $r \geqslant 5$, near criticality $\mathscr{G}_{n,p}$ has no K_r-percolating subgraphs of order $1 \ll k \leqslant n^c$, for some $0 < c < 1$; see Section 2.1.

We thus take a different approach for finding the sharp threshold for the K_r-dynamics. Rather than studying K_r-percolating subgraphs of $\mathscr{G}_{n,p}$, we carry out a systematic study of its *witness graphs*, that is, inclusion-minimal subgraphs of $\mathscr{G}_{n,p}$ that add a given edge. The parameter λ appears in the exponent of $p_c(n,K_r)$ because any witness graph for a given edge with v additional vertices contains at least $\lambda v + 1$ edges. The constant γ that appears in $p_c(n,K_r)$ is connected to the Fuss–Catalan numbers since the graphs that most efficiently add a given edge, which we will call *tree witness graphs (TWGs)*, have a certain $\left(\binom{r}{2}-1\right)$-ary tree-like structure; see Figure 2 below.

Establishing sufficiently precise bounds for the number of witness graphs that are more "costly" than TWGs turns out to be highly non-trivial, and much of the current work is devoted to this task. The central challenge, in analyzing the structure of such graphs, is to control the spread of "zero-cost" edges that occurs after a "costly" (non-tree-like) step in the K_r-dynamics. We refer the reader to Section 2 for a more in-depth overview of the challenges and our main ideas in this work.

In addition, our proof of Theorem 1.1 shows that in the subcritical regime almost all edges added to $\mathscr{G}_{n,p}$ arise from TWGs. This allows us to derive the following sharp estimate for the number of edges in $\langle \mathscr{G}_{n,p} \rangle_{K_r}$.

Theorem 1.2. *Fix $r \geqslant 5$. Let $p = (\bar{\gamma} n)^{-1/\lambda}$ for some $\bar{\gamma} > \gamma$, where γ is as in* (1.3) *above. Define*

$$\bar{\alpha} := (r-2)!\bar{\gamma}^{r-2} > \alpha_{\binom{r}{2}-2},$$

and let $1 < \rho < 1 + \left(\binom{r}{2}-2\right)^{-1}$ be the smallest root of the equation

$$\rho^{\binom{r}{2}-1} = \bar{\alpha}(\rho - 1).$$

Then,

$$|E(\langle \mathscr{G}_{n,p} \rangle_{K_r})| \sim \rho \cdot p \binom{n}{2},$$

in probability, as $n \to \infty$.

In other words, for subcritical p, the number of edges in $\mathscr{G}_{n,p}$ essentially only increases by a factor of ρ, universally bounded by $1 + \left(\binom{r}{2}-2\right)^{-1}$, after the K_r-dynamics have stabilized. See Figure 1 for an illustration of our main results (however, see also Section 8.2 for a discussion on the limitations of such simulations).

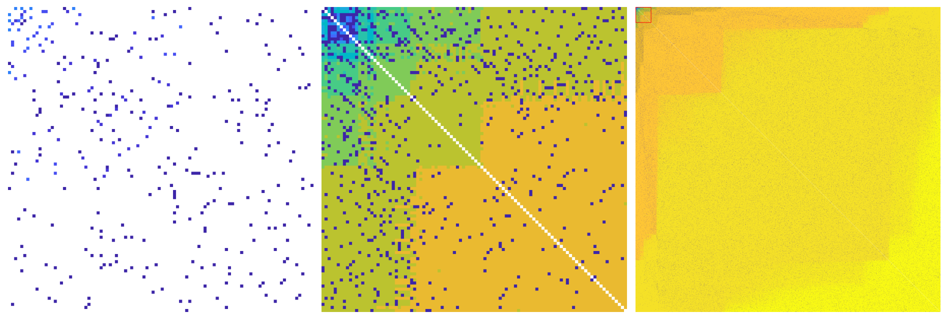

(A) (B) (C)

FIGURE 1. Simulations of the K_5-dynamics on random graphs with $n = 2000$ vertices, before and after criticality. Each point (i, j) is colored along a gradient from dark blue to yellow according to the time at which the edge $\{v_i, v_j\}$ is added in the K_5-dynamics (obtaining a symmetric image), where the vertices are ordered retrospectively, biased so that vertices incident to many early-added edges appear earlier. Dark blue points represent edges present in the initial graph, yellow those added in the final round, and white the edges that are never added. (a) In the sub-critical dynamics, the edge density only increases by a constant factor (few lighter blue pixels appear). The image shows 100 vertices, as no percolation occurs in other vertices. The process finished in 4 rounds. (b) The early stages of the super-critical dynamics (dark blue to green) on the first 100 vertices in the biased order. (c) The process "explodes," and all remaining edges are added (orange to yellow). The center image is the upper left region (highlighted in red) in the right image. The process finished in 15 rounds.

1.2. **Related work.** There is a large mathematical literature on bootstrap percolation. In particular, the r-neighbor dynamics on Euclidean lattices $\mathbb{Z}^d$ and cubes $[n]^d$ is well studied; see, e.g., [1, 8, 19, 31, 34, 46, 52]. Moreover, Balister, Bollobás, Morris and Smith [7] (and see references therein) have recently proven the *universality conjecture* concerning more general forms of monotone cellular automata. We also note that Janson, Łuczak, Turova and Vallier [35] (cf. [4, 26, 49]) have studied the r-neighbor dynamics on $\mathscr{G}_{n,p}$ in great detail.

In graph theory, there has been continued interest in weak saturation and graph bootstrap percolation since the works of Bollobás [17] and Balogh, Bollobás, and Morris [9]. Bollobás [17] initiated the study of the minimal number of edges in a weakly K_r-saturated subgraph of K_n, identifying this number when $3 \leqslant r \leqslant 6$. Alon [2], Frankl [27], Kalai [36, 37] and Lovász [41] (cf. Yu [54]) generalized the result to all r. Later, Balogh, Bollobás, Morris and Riordan [10] obtained a hypergraph analogue of this classical result. See also [6,15,24,37,40] (and references therein) for the bipartite case, and [47, 48, 50] for the asymptotic behavior in the general case. Many other related problems have been studied. For example, there are several works on the running time (number of rounds until the H-dynamics stabilize); see, e.g., [11, 18, 23, 25, 29, 30, 44]. Another line of research, initiated by Korándi and Sudakov [39], studies the minimal number of edges in a subgraph of $G \sim \mathscr{G}_{n,p}$ such that the H-dynamics, restricted to the "host graph" G, eventually adds all the edges in G; see also [5, 16, 21] for more recent progress in this direction.

As discussed above, Balogh, Bollobás, and Morris [9] initiated the study of threshold probabilities for H-percolation, primarily focusing on the case that $H = K_r$ is a clique. Some other types of template graphs H have been studied in [9] and elsewhere, e.g, in [12–15]; however, a general understanding for *all* H (see Problem 1 in [9]) remains widely open.

In closing, let us mention that a sharp Fuss–Catalan-type threshold, resembling Theorem 1.1, was recently discovered in the context of hypergraph bootstrap percolation [42]. In that work, the challenge in analyzing the counterparts of witness graphs was tackled by exterior algebraic shifting. We tried to apply similar methods in this work, via rigidity theory, but we were unsuccessful due to some fundamental obstructions; see Section 8.4 below. On the other hand, it seems likely to us that the strategies presented in the current work are robust, and will be capable of giving a "combinatorial" proof of the results in [42], and more importantly, useful in studying various general versions of graph and hypergraph bootstrap percolation, where a direct analysis of the dynamics is typically challenging.

1.3. **Outline.** In Section 2 we give a detailed overview of our proof strategy. Afterward, in Section 3 we prove the upper bound in Theorem 1.1, and in Sections 4–6 we prove the lower bound. In Section 4 we set up terminology for the various types of steps that can be taken in the construction of general witness graphs, and prove several fundamental properties of such graphs that will be used in the proof of the lower bound. Section 5 establishes a coarse lower bound $p_c = \Theta(n^{-1/\lambda})$, which already solves Problem 3 in [9]. Building on these arguments, and adding several more new ideas, we find the sharp lower bound in Section 6. Finally, in Section 7 we prove Theorem 1.2 on the addition of edges in the subcritical setting.

CONTENTS

2. CONTEXT AND PROOF OVERVIEW

The purpose of this section is to explain why and in what ways the K_r-dynamics for $r \geqslant 5$ is different and more complex than when $r = 4$, and to outline the main ideas behind our proofs in this setting.

2.1. **Nucleation in the K_r-dynamics.** As observed in [9, §4], any graph that is K_4-percolating can be decomposed into two or three percolating subgraphs. Therefore, K_4-percolating graphs have percolating subgraphs of all orders. A similar property, for classical bootstrap percolation on Euclidean lattices, was first observed by Aizenman and Lebowitz [1], and analogues of this property have since been used in various other contexts.

The simplest K_4-percolating graphs, that are grown recursively, by "nucleation," were already recognized in [17, Fig. 4]. As discussed in Section 1, to construct such a graph, we start with a single edge and then iteratively add vertices to the graph with two neighbors amongst those previously added. The growth of such a graph in $\mathscr{G}_{n,p}$, started from one of its edges, gives rise to time-inhomogeneous branching process with an increasing birthrate (since, as this graph grows, it becomes more likely to find a vertex outside of it with at least two neighbors inside). When $p = (\alpha n \log n)^{-1/2}$, survival to time $\alpha \log n$ creates a "critical droplet" (the branching process become critical), after which the full percolation of $\mathscr{G}_{n,p}$ is all but inevitable. As it turns out, for $\alpha < 3$, with high probability, the branching process started at *some* edge of $\mathscr{G}_{n,p}$ becomes super-critical. All other ways of percolating are of smaller order, and the sharp threshold occurs at $p_c \sim (3n \log n)^{-1/2}$ [3,4,9,38]. In other words, the most likely way for $\mathscr{G}_{n,p}$ to K_4-percolate is to have an edge $e \in E(\mathscr{G}_{n,p})$ whose vertices form a *contagious set* (i.e., a set whose initial infection leads to full infection) for the classical 2-neighbor bootstrap percolation process on $\mathscr{G}_{n,p}$.

For $r \geqslant 5$, one can similarly construct K_r-percolating graphs via "nucleation" by starting with a K_{r-2} and then iteratively adding vertices connected to $r-2$ previously added vertices. In fact, the series of works [2, 17, 27, 36, 37, 41] have established that every k-vertex K_r-percolating graph has at least $(r-2)k - \binom{r-1}{2}$ edges, precisely the same number of edges as in these "nucleating" constructions.

One might expect that the growth of such a graph in $\mathscr{G}_{n,p}$ is the cause for K_r-percolation. However, a simple union bound implies that, when p is near $p_c = n^{-1/\lambda+o(1)}$, the graph $\mathscr{G}_{n,p}$ typically contains no percolating subgraphs with k vertices in the range $1 \ll k \ll L$, where

$$L = (np^{r-2})^{-1/(r-3)} = n^{(r-4)/(r^2-r-4)+o(1)}. \tag{2.1}$$

This indicates a qualitative change occurring at $r \geqslant 5$, as $\lambda = r-2$ when $r = 4$, but $\lambda < r-2$ for $r \geqslant 5$. Consequently, the Aizenman–Lebowitz-type property fails for K_r-percolation with $r \geqslant 5$; that is, there exist K_r-percolating graphs that do not contain percolating subgraphs of all smaller orders. This helps explain why K_r-percolating graphs are generally much less understood when $r \geqslant 5$. In addition, (2.1) says that if a "critical droplet" for K_r-percolation in $\mathscr{G}_{n,p}$ does form, it must appear rather abruptly (see Section 8.1). Our proof does not identify such a droplet directly; instead, it takes a different approach based on the analysis of *witness graphs*.

2.2. **The coarse bounds for p_c.** The concept of a witness graph plays a key role in [9], where the coarse bounds (1.1) are proved. A witness graph W for an edge e is an inclusion-minimal subgraph for which $e \in E(\langle W \rangle_{K_r})$; see Definiton 4.2 below. The existence of such a witness graph W as a subgraph of G bears witness to the fact that the K_r-dynamics starting from G eventually add e to the graph. As observed in [9], witness graphs satisfy an Aizenman–Lebowitz-type property. Specifically (see Lemma 4.4 below), in each step of the K_r-dynamics, the maximal number of edges in a witness graph, over all witness graphs for edges added so far, can increase by at most a factor of $\binom{r}{2}$, the number of edges in K_r.

The bounds (1.1) are obtained as follows in [9]. To obtain the upper bound in (1.1), the authors first show that, when p is sufficiently supercritical, most edges in K_n are added by a *K_r-ladder graph* (see Figure 2 below) of logarithmic height, and then full percolation is achieved by sprinkling. Informally, such a graph is obtained by stringing together a series of copies of K_r, with its target edge e in a copy at one of the ends, and then removing all shared edges and e. The lower bound in (1.1), on the other hand, is obtained by first showing that all witness graphs with k vertices have at least as many edges $\lambda(k-2)+1$ (see also [12, Lemma 12] for an alternative proof of this fact, that extends to all graphs H) as there are in K_r-ladder graphs of the same size, and then applying the Aizenman–Lebowitz property for witness graphs.

2.3. **Sharp upper bound on p_c.** A limitation of the approach in [9] is that, apart from determining the minimal number of edges, it does not examine in detail the combinatorial structure of witness graphs. The first step in finding a sharp upper bound on p_c is to identify the class of all edge-minimal witness graphs. As it turns out, these are $(\binom{r}{2}-1)$-ary tree-like graphs that we call *K_r-tree witness graphs (TWGs)*, and this explains the "Fuss–Catalan-related" constant γ in Theorem 1.1. Informally, a TWG for an edge e is obtained by first stringing together copies of K_r in a tree-like manner, so that e is in one (root) copy, and every edge is in at most two. Then, the TWG is obtained by removing e and all such shared edges from the graph. Note that these removed edges, and e in particular, will be added back by the K_r-dynamics. See Figure 2 for examples.

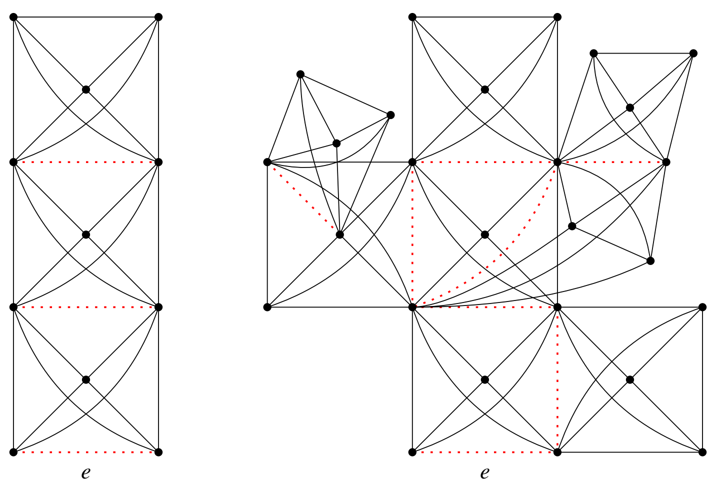

FIGURE 2. A K_5-ladder (at left) of height 3 and a more complicated K_5-tree (at right) of order 8. The non-edges of these K_5-trees that get added by the K_5-dynamics are indicated by red dotted lines. Their target edges e are labelled.

To establish the upper bound in Theorem 1.1 we show (see Section 3 below) that, in the supercritical regime, with high probability, all edges are eventually added to $\mathscr{G}_{n,p}$ by the K_r-dynamics using a TWG of logarithmic "order" (number of copies of K_r used in its construction). This is obtained by studying the expectation and variance of the number of TWGs for a given edge e in $\mathscr{G}_{n,p}$. The central idea is that every subgraph S of a TWG T for some e has strictly smaller edge density than T, and moreover, the closer the density of S is to that of T, the more closely its structure resembles a TWG for e. The technicalities, however, are somewhat involved.

2.4. **Sharp lower bound on p_c.** The sharp lower bound in Theorem 1.1 is more challenging, and it is here that much of the innovation in this work occurs. We obtain it by excluding the appearance of *all* types of witness graphs for a given edge e. As mentioned above, TWGs are the only edge-minimal witness graphs, where a k-vertex TWG has $\lambda(k-2)+1$ edges. Overall, our strategy is to show that a witness graph with k vertices and $\lambda(k-2)+1+\chi$ many edges is, in some useful sense, $O(\chi)$-close to being a TWG. We call χ its *excess* number of edges.

The *red edge algorithm (REA)* introduced in [9] is a deterministic procedure that constructs a witness graph W for a target edge e. At each step, the edges added so far form a collection of edge-disjoint components, and the algorithm introduces a new copy of K_r that merges all components sharing edges with it. In the final step, the copy of K_r that is added includes the target edge e and it merges all existing components into a single one.

There are three types of steps in the REA. First, *tree steps*, in which components are merged in a tree-like manner. For example, a witness graph is a TWG precisely when all its REA steps are tree steps. Second, *costly steps*, which increase the excess χ of the witness graph, typically by merging components in a non–tree-like way. Third, *zero-cost steps*, which take place within a single component and do not increase the excess. While costly steps themselves are harmless (since their impact is accounted for by the excess χ), they may initiate cascades of zero-cost steps, potentially giving rise to a highly complicated structure within the witness graph. Bounding the spread of zero-cost steps in terms of the excess is far from trivial, and to address this problem, we introduce two key new ideas.

2.4.1. *The tree decomposition.* In Section 5, we introduce the *tree decomposition* of a witness graph, and describe how it is constructed over time as the REA unfolds. In each step, we decompose all existing components into a "bad part" and various other "tree parts." These tree parts are similar to TWGs, but more general, as they may have multiple "target" edges that have yet to be reused by the REA. We do this carefully enough so that, on the one hand, any zero-cost spread of edges occurs only within the bad part, and so that, on the other hand, the number of inclusion-maximal tree parts that appear during the REA construction is bounded by $O(\chi)$.

We show that this decomposition is, on its own, enough to prove that $p_c = \Theta(n^{-1/\lambda})$, which already answers Problem 3 in [9]. We note that the assumption that $r \geqslant 5$ is heavily used here. Indeed, when $r=4$, a TWG itself creates a large zero-cost spread of edges (in fact, it turns into a clique), so it is impossible to distinguish between the bad part and the tree parts of a component. Indeed, when $r=4$, the lower bound $\Omega(n^{-1/\lambda})$ fails, recalling that $p_c(n,K_4) \sim (3n\log n)^{-1/2}$.

2.4.2. *The $(r-2)_*$-bootstrap percolation process.* In Section 6, to establish the sharp lower bound in Theorem 1.1, we refine our control on the spread of zero-cost edges within the various tree parts of the tree decomposition. Using this refinement, we deduce that the tree parts are all close to being TWGs, and the sharp lower bound follows.

It is not hard to see that, when considered on its own, a tree part T does not permit the propagation of zero-cost edges. However, if some of its vertices are involved in costly steps of the REA, this may cause additional edges to spread throughout T. Controlling this spread is the main objective in Section 6.

To this end, we make a comparison with a certain variation of the $(r-2)$-neighbor dynamics on the vertices of T, which we call the *$(r-2)_*$-bootstrap percolation process*. The "seeds" (i.e., initially infected vertices) in this process are the vertices of T that are involved in costly steps. This process evolves like the usual $(r-2)$-dynamics, except for one special rule. Being a tree part, T consists of copies of K_r arranged so that each new copy intersects the union of the preceding ones in exactly one edge. In the modified dynamics, we allow a vertex v which belongs to such a shared edge f to be infected at a lower threshold of $r-3$ if it has $r-4$ infected neighbors in one copy of K_r containing f, and another infected neighbor in a different such copy, assuming that none of these neighbors is the other endpoint of f.

The close connection between the edge K_r-dynamics and the vertex $(r-2)_*$-dynamics follows from the observation that the propagation of new edges in the K_r-dynamics can "pass through" a shared edge f, from one part of T to another only if there exists a vertex $v \in f$ satisfying the above condition. Otherwise, the "hinge" (i.e., shared edge) f in such a T acts as a sort of "bottleneck" for the spread of new edges throughout T. See Section 6.1 and Figure 3 below for more details.

Finally, even though the tree parts in our decomposition are not trees in the usual sense, they possess a workable tree-like structure. This allows us to analyze the $(r-2)_*$-bootstrap process on a tree part by extending the ideas in the work of Riedl [45] on the r-neighbor dynamics in trees.

3. Sharp upper bound

In this section, we prove the sharp upper bound in Theorem 1.1.

Proposition 3.1 (Sharp upper bound)**.** *Fix $r \geqslant 5$ and let $\varepsilon > 0$. Put*

$$p = \left(\frac{1+\varepsilon}{\gamma n}\right)^{1/\lambda}.$$

Then, with high probability, $\langle \mathscr{G}_{n,p} \rangle_{K_r} = K_n$.

Throughout this section, we will use c to denote a constant depending only on r. The specific value of c may change, but it will always represent a constant depending only on r. We may also use c', etc., sometimes to emphasize that its value is changing from one line to another.

3.1. **Growing TWGs.** Let us now formalize the notion of a TWG, as discussed briefly in Section 2.3 above. In fact, we will give two recursive definitions of such structures, depending on whether it is more convenient to, in a certain sense (to be made precise below), grow the TWG from its *root* or by *grafting a leaf* somewhere along one of its *branches*.

Definition 3.2 (Growing a TWG from the root)**.** A graph T is called a *tree witness graph (TWG)* for the edge e if either

(1) $T = e$; or else
(2) for some copy $H^{(e)}$ of K_r with $e \in E(H^{(e)})$ there are TWGs T_f for each edge $f \in E(H^{(e)} \setminus e)$ such that T_f and $T_{f'}$ are vertex disjoint outside of $H^{(e)}$ for all $f \neq f' \in E(H^{(e)} \setminus e)$, and $T = \bigcup_{f \in E(H^{(e)} \setminus e)} T_f$.

In case (2) we call $H^{(e)}$ the *root* of T and the T_f its *primal branches*.

Next, we define the order of a TWG. Informally, this is the number of copies of K_r in its associated hyper-tree.

Definition 3.3 (Order of TWGs)**.** The *order* $\vartheta(T)$ of a TWG T for e is defined inductively. If T consists of a single edge then $\vartheta(T) = 0$, otherwise

$$\vartheta(T) = 1 + \sum_{f \in E(H^{(e)} \setminus e)} \vartheta(T_f),$$

where $H^{(e)}$ is the root of T. If $\vartheta(T) = k$, then we call T a *k-TWG*.

Definition 3.4 (Branches and leaves of TWGs)**.** Let T be a TWG with root $H^{(e)}$. We call $B \subset T$ a *branch* of T if B is a TWG. Inductively, B is a branch of T if $B = T$ or if B is a branch of one of its primal branches T_f. A branch L of order $\vartheta(L) = 1$ is called a *leaf* of T.

In particular, note that if B is a branch of a TWG T, then B is a TWG for some edge $e' \notin E(T)$.

It will sometimes be more convenient to construct larger TWGs by removing an edge from a smaller TWG and replacing it with a leaf. Hence, we also give the following equivalent definition of TWGs.

Definition 3.5 (Grafting a leaf to a TWG)**.** A graph T is a 0-TWG for an edge e if $T = e$. For $k \geqslant 1$, a graph T is a k-TWG for e if there exists

- a $(k-1)$-TWG T' for e;
- an edge $e' \in E(T')$; and

- a copy H' of K_r containing e', and no other vertices in T' beside the endpoints of e',

such that T is obtained from T' by

- removing the edge e'; and then
- *grafting* (taking the union with) the leaf L, obtained from H' by removing e', in its place.

Finally, we note that the following facts can be proved by a simple induction, using either of the above equivalent definitions.

Lemma 3.6. *Suppose that T is a k-TWG for an edge e. Then*

- $v(T) = (r-2)k+2$;
- $e(T) = \lambda(r-2)k+1$; *and*
- $e \in E(\langle T\rangle_{K_r})$.

In other words, if a graph G contains a TWG T for an edge e, then T is indeed *witness* to the fact that $e \in E(\langle G\rangle_{K_r})$ is eventually added by the K_r-dynamics. For the other two facts, simply note that T contains the edge e and its two vertices, and that, each time a TWG is grown by adding a leaf, $r-2$ vertices and $\lambda(r-2) = \binom{r}{2} - 2$ edges are added.

3.2. **Counting TWGs.** We recall that the *Fuss–Catalan number*

$$\mathrm{FC}_d(k) = \frac{1}{dk+1}\binom{(d+1)k}{k} \tag{3.1}$$

is the number of $(d+1)$-ary trees with k internal nodes and $dk+1$ leaves. When $d=1$, $\mathrm{FC}_1(k)$ is the well-known *Catalan number* $C(k)$.

By Stirling's approximation, the Fuss–Catalan numbers have asymptotic growth rate

$$\alpha_d = \frac{(d+1)^{d+1}}{d^d}.$$

More specifically, for any fixed d, it follows by (3.1) that

$$\mathrm{FC}_d(k) \sim \beta_d \frac{\alpha_d^k}{k^{3/2}}, \tag{3.2}$$

as $k \to \infty$, where

$$\beta_d = \sqrt{\frac{d+1}{2\pi d^3}}.$$

Finally, we note that, similarly to the Catalan numbers, the Fuss–Catalan numbers can be defined recursively by setting $\mathrm{FC}_d(0) = 1$ and, for $k \geqslant 1$,

$$\mathrm{FC}_d(k) = \sum_{k_1+\cdots+k_{d+1}=k-1} \prod_{i=1}^{d+1} \mathrm{FC}_d(k_i). \tag{3.3}$$

Definition 3.7. Let $t(k)$ denote the number of (labelled) k-TWGs for a given edge e on a given vertex set of size $(r-2)k+2$.

Lemma 3.8 (Number of TWGs)**.** *For all integers $k \geqslant 0$, we have that*

$$t(k) = \frac{((r-2)k)!}{(r-2)!^k} \mathrm{FC}_{\binom{r}{2}-2}(k). \tag{3.4}$$

Note that, by (1.3), (3.2) and (3.4), it follows that

$$t(k) \sim \beta_{\binom{r}{2}-2} \frac{\gamma^{(r-2)k}}{k^{3/2}} ((r-2)k)!, \tag{3.5}$$

as $k \to \infty$.

Proof. The proof is by induction on k, combining Definition 3.2 with the recurrence (3.3).

The base case $k=0$ is trivial, since $t(0)=1$. On the other hand, for $k>0$, to construct a k-TWG T for e, we need to choose

(1) the other $r-2$ vertices in the root $H^{(e)}$ of T;
(2) the orders $k_f = \vartheta(T_f)$ for each of its $\binom{r}{2}-1$ many primal branches T_f, in such a way that $\sum_f k_f = k-1$;
(3) a partition $(\Pi_f)_f$ of the remaining $(r-2)(k-1)$ vertices outside of $H^{(e)}$, in such a way that each $|\Pi_f| = (r-2)k_f$; and
(4) a k_f-TWG for each primal branch T_f.

As such, it follows that

$$t(k) = \binom{(r-2)k}{r-2} \sum_{\sum_f k_f = k-1} ((r-2)(k-1))! \prod_f \frac{t(k_f)}{((r-2)k_f)!}.$$

Consequently, the sequence

$$\frac{(r-2)!^k}{((r-2)k)!} t(k)$$

satisfies (3.3) with $d = \binom{r}{2}-1$, and the result follows. ■

3.3. **Balanced TWGs.** The following concept is needed for technical reasons that will appear later in the proof of Proposition 3.1.

Definition 3.9 (Balanced TWGs)**.** Let T be a TWG for an edge e, with root $H^{(e)}$, and primal branches T_f, as in Definition 3.2. We say that T is *balanced* if the orders $\vartheta(T_f)$ are all equal.

Note that this definition is *not* recursive. That is, crucially, the primal branches T_f themselves need not be balanced.

Definition 3.10. Let $b(k)$ denote the number of (labelled) balanced k-TWGs for a given edge e on a given vertex set of size $(r-2)k+2$.

Note that, if T is balanced then $k-1$ is divisible by $\binom{r}{2}-1$, in which case we put

$$q = \frac{k-1}{\binom{r}{2}-1}, \tag{3.6}$$

so that all of its primal branches are of order $\vartheta(T_f) = q$.

Lemma 3.11 (Number of balanced TWGs)**.** *For all integers $k \geqslant 0$ such that $k-1$ is divisible by $\binom{r}{2}-1$, we have that*

$$b(k) = \frac{((r-2)k)!}{(r-2)!}\left(\frac{t(q)}{((r-2)q)!}\right)^{\binom{r}{2}-1}. \tag{3.7}$$

Proof. To construct a balanced k-TWG T for a given edge e on a given vertex set of size $(r-2)k+2$, we need to choose

(1) the $r-2$ other vertices in the root $H^{(e)}$;
(2) a partition of the $(r-2)(k-1)$ remaining vertices into $\binom{r}{2}-1$ equal sets of size $(r-2)q$; and
(3) $\binom{r}{2}-1$ TWGs of order q for the primal branches.

Hence

$$b(k) = \frac{((r-2)k)!}{(r-2)!((r-2)q)!^{\binom{r}{2}-1}} t(q)^{\binom{r}{2}-1},$$

giving the result. ■

The key point is that the difference between $b(k)$ and $t(k)$ is not seen at the exponential scale. Indeed, by (3.5) and (3.7) it follows that

$$b(k) \sim \phi \frac{t(k)}{k^{(3/2)(\binom{r}{2}-2)}}, \tag{3.8}$$

where

$$\phi = \frac{\left(\binom{r}{2}-1\right)^{3/2} \beta_{\binom{r}{2}-2}^{\binom{r}{2}-2}}{(r-2)!\gamma^{r-2}}$$

is a constant depending only on r.

Finally, the last result in this section implies that in the supercritical setting, we have, in expectation, many balanced k-TWGs for any given edge e for $k = \beta \log n$, provided that β is sufficiently large.

Lemma 3.12. *Fix $e \in E(K_n)$. Let $\varepsilon > 0$ and put*

$$p = \left(\frac{1+\varepsilon}{\gamma n}\right)^{1/\lambda}.$$

Suppose that $k = k(n)$ is a sequence such that $k-1$ is divisible by $\binom{r}{2}-1$ for all n, and that $1 \ll k \ll \sqrt{n}$, as $n \to \infty$. Then, as $n \to \infty$, the number $B_k^{(e)}$ of balanced k-TWGs for e in $\mathscr{G}_{n,p}$ satisfies

$$\mathbb{E}(B_k^{(e)}) \sim \phi' \frac{(1+\varepsilon)^{(r-2)k}}{k^{(3/2)(\binom{r}{2}-1)}} p,$$

where

$$\phi' = \beta_{\binom{r}{2}-2} \phi.$$

Proof. By linearity of expectation,

$$\mathbb{E}(B_k^{(e)}) = \binom{n-2}{(r-2)k} b(k) p^{\lambda(r-2)k+1}.$$

Since $1 \ll k \ll \sqrt{n}$, we have that

$$\binom{n-2}{(r-2)k} \sim \frac{n^{(r-2)k}}{((r-2)k)!}.$$

Therefore, by (3.5) and (3.8), we find that

$$\mathbb{E}(B_k^{(e)}) \sim \beta_{\binom{r}{2}-2} \phi \frac{(\gamma n p^{\lambda})^{(r-2)k}}{k^{(3/2)(\binom{r}{2}-1)}} p,$$

and the result follows. ■

3.4. **Partial TWGs.** The next key step towards the proof of Proposition 3.1 is studying subgraphs of TWGs.

Definition 3.13 (Partial TWGs)**.** We call a subgraph S a *partial TWG* for an edge e if there exists a TWG T for e such that $S \subsetneq T$ is a proper edge-induced subgraph of T.

We will need to quantify, at least approximately, how far a partial TWG S is from being a complete TWG. In the case that S is the intersection of two potential TWGs T_1 and T_2, this information will help bound the correlation between the events that T_1 and T_2 are present in $\mathscr{G}_{n,p}$.

For a graph G and some $v \in V(G)$, we let $E_v(G) \subset E(G)$ denote the edges of G that are incident to v.

Definition 3.14 (Sets F, P and D)**.** Let T be a TWG for an edge e. Suppose that $S \subsetneq T$ is a partial TWG for e. We partition the set of vertices in T that are not in e into the sets F, P, and D of vertices v for which $E_v(T)$ is *fully* contained in, *partially* contained in, and *disjoint* from $E(S)$.

In other words, for each such $v \notin e$, we put

- $v \in F$ if $E_v(T) \subset E(S)$;
- $v \in D$ if $E_v(T) \cap E(S) = \emptyset$; and

- $v \in P$ otherwise.

Note that that F, P and D are pairwise disjoint since S is edge induced and thereby has no isolated vertices.

To be precise, we can write $F = F(S,T)$, $P = P(S,T)$ and $D = D(S,T)$ as all three sets depend on S and T. However, the set $F \cup P = V(S) \setminus e$ depends only on S, and we denote

$$\sigma(S) = |V(S) \setminus e|$$

the number of vertices in S that are not in e.

Finally, let us introduce the three key parameters associated with $S \subsetneq T$ which will help us measure how far S is from being a TWG.

Definition 3.15 (Key partial TWG parameters)**.** We define the

- *edge efficiency* of S as $\mathscr{E}(S) = \lambda\sigma(S) - e(S)$;
- *within tree deficiency* of S as $\mathscr{D}(S,T) = |P|$; and
- *tree extendability* $\mathscr{T}(S)$ of S to be the minimum of $\sigma(T_*) - \sigma(S)$ over all TWGs T_* for e satisfying $S \subsetneq T_*$.

As indicated by the notation, $\mathscr{D}$ depends on S and T, but $\mathscr{E}$ and $\mathscr{T}$ are, in fact, independent of T.

3.5. **Structure of partial TWGs.** The goal of this subsection is to establish the following lemma that connects between the three key partial TWG parameters.

Lemma 3.16 (Partial TWG structural lemma)**.** *Let T be a TWG and $S \subsetneq T$ a partial TWG. Then,*

(1) $\mathscr{E}(S) \geqslant 0$, with equality if and only S is obtained from T by removing one of its branches;
(2) $\mathscr{D}(S,T) \leqslant c\mathscr{E}(S) + 2$; and
(3) $\mathscr{T}(S) \leqslant c\mathscr{E}(S)$,

for some constant c depending only on r.

Recall that a branch $B \subset T$ is a TWG for some $e' \notin E(T)$. When *removing* B from T, we do *not*, however, mean to remove the two vertices in e' from what remains.

The proofs of these results, in Sections 3.5.3, 3.5.4 and 3.5.5 below, are by a somewhat technical case analysis. The reader might wish to first skip ahead to Sections 3.6 and 3.7 to see how it all fits together, before returning to the proof of this lemma.

Parts (1) and (2) are closely related, and the proof of (2) is essentially a continuation of the arguments for (1). Part (3), on the other hand, requires a different type of case analysis.

3.5.1. *Strictly-balanced graphs.* The proof of Lemma 3.16 makes critical use of the fact that the clique K_r is strictly balanced for $r \geqslant 5$. To highlight this fact, we present Lemma 3.19 below in the broader setting of strictly balanced graphs; Lemma 3.20, concerning the clique K_r for $r \geqslant 5$, will follow as a special case useful in our arguments.

Definition 3.17 (Strictly balanced)**.** A graph H is *strictly balanced* if

$$\frac{e(F)-1}{v(F)-2} < \lambda(H), \tag{3.9}$$

for every subgraph $F \subset H$ with $3 \leqslant v(F) < v(H)$, where

$$\lambda(H) = \frac{e(H)-2}{v(H)-2} \tag{3.10}$$

is its λ*-density* (extending (1.2) above).

Definition 3.18. For a strictly-balanced graph H, we let $\xi(H) > 0$ denote the minimal possible value of

$$\lambda(H)(v(F)-2)-(e(F)-1)$$

over all $F \subset H$ with $3 \leqslant v(F) < v(H)$.

Lemma 3.19. *Suppose that H is a strictly balanced graph and that K is a proper subgraph of the graph $H \setminus e$, obtained from H by removing some $e \in E(H)$. Let x denote the number of vertices in K that are not in e. Then we have that $0 \leqslant x \leqslant v(H)-2$ and*

$$e(K) \leqslant \lambda(H)x,$$

with equality if and only if either

(1) $x = 0$ and so $e(K) = 0$; or else

(2) $x = v(H)-2$ and $e(K) = e(H)-2$.

In all other cases, in fact,

$$e(K) \leqslant \lambda(H)x - \xi(H),$$

where $\xi(H) > 0$ is as in Definition 3.18.

Proof. First, we note that if K contains $x = 0$ vertices outside of e, then clearly $e(K) = 0$ since $K \subset H \setminus e$, and so $\lambda(H)x - e(K) = 0$ holds trivially.

Next, if $x = v(H)-2$, then

$$\lambda(H)x - e(K) = e(H)-2-e(K) \geqslant 0$$

by (3.10) and since $K \subsetneq H \setminus e$. Moreover, equality holds if and only if $e(K) = e(H)-2$.

Finally, suppose that $0 < x < v(H) - 2$. In this case, consider $F = K \cup e$ obtained from K by adding e. Observe that

$$3 \leqslant v(F) = x + 2 < v(H).$$

Then, since H is strictly balanced, it follows by (3.9) that

$$e(K) = e(F) - 1 \leqslant \lambda(H)(v(F) - 2) - \xi(H) = \lambda(H)x - \xi(H),$$

as claimed. ■

3.5.2. *Proof of Lemma 3.16.* Let us return to the setting of Lemma 3.16, where $S \subsetneq T$ is a proper edge-induced subgraph of a k-TWG T, for some $k \geqslant 1$. By Definition 3.5, T can be obtained by removing an edge e' from some $(k-1)$-TWG T' and then grafting a leaf L in its place. Therefore, since $L \cup e'$ is a copy of K_r, the next useful result follows immediately by Lemma 3.19.

Lemma 3.20 (Efficiency within a leaf)**.** *Suppose that L is not contained in S. Let x denote the number of vertices in $S \cap L$ that are not in e'. Then*

$$e(S \cap L) \leqslant \lambda x,$$

with equality if and only if

(1) $x = 0$*; or else*
(2) $x = r - 2$ *and* $e(S \cap L) = \binom{r}{2} - 2$*,*

and in all other cases,

$$e(S \cap L) \leqslant \lambda x - \xi,$$

for some $\xi > 0$ *depending only on* r*.*

In our proofs, it will often be useful to consider one of the two following partial TWGs of T'.

- Let S' be the edge-induced subgraph of T', induced by the edges of S that are contained in T'.
- Let $\bar{S}'$ be obtained from S' by adding the edge e' and its vertices (if not already included in S').

Note that because $S \subset T$ and $e' \notin E(T)$, it follows that e' is not an edge in S'. Therefore, $S' \subsetneq T'$, which will allow us to apply various induction hypotheses directly in some cases. On the other hand, $\bar{S}'$ may be equal to T', in which case other arguments will be required.

3.5.3. *Efficiency: Proof of Lemma 3.16(1).* First, we observe that if S is obtained from T by removing a branch B that is a TWG for some $f \notin E(T)$ then $T_f = S \cup f$ is a k'-TWG for e for some $k' \geqslant 0$. In this case, $\sigma(S) = v(T_f) - 2 = k'(r-2)$ and $e(S) = e(T_f) - 1 = (\binom{r}{2} - 2)k'$, and so $\mathscr{E}(S) = 0$.

Next, for the inequality and the other direction of the equality case, we proceed by induction on k, the order of T.

The base case $k = 0$ holds trivially, as then T is a single edge and $S = \emptyset$, which can be obtained from T by removing T (a branch of T) from T.

For the induction step, suppose that $k \geqslant 1$. Then, as discussed above, T can be obtained by removing an edge e' from some $(k-1)$-TWG T' and then grafting a leaf L in its place.

Empty case. Suppose that S contains no edges in L and no vertices in L, except possibly those in e'. Then $S = S'$. Moreover, we have that $S \subsetneq T'$, since $e' \in E(T')$ but $e' \notin E(T)$. Therefore, the inductive hypothesis applies directly to show that $\mathscr{E}(S) = \mathscr{E}(S') \geqslant 0$. For the case of equality, note that if S is obtained from T' by removing a branch B' from T', then $e' \in E(B')$. In this case, S can be obtained from T by removing the branch B obtained from B' by removing e' and grafting L in its place.

Complete case. Next, suppose that $L \subset S$. Then $\bar{S}' \subsetneq T'$ since $S \subsetneq T$. Hence by induction, we have that $\mathscr{E}(\bar{S}') \geqslant 0$, and it follows that $\mathscr{E}(S) \geqslant 0$ since S has precisely $r-2$ more vertices and $\binom{r}{2} - 2 = \lambda(r-2)$ more edges than $\bar{S}'$, so that $\mathscr{E}(S) = \mathscr{E}(\bar{S}') \geqslant 0$. For the case of equality, recall that $e' \in E(\bar{S}')$. Therefore, if $\bar{S}'$ is obtained from T' by removing one of its branches B, then $e' \notin E(B)$, in which case B is also a branch of T, and S can be obtained by removing B from T.

Nearly-complete case. Next, we suppose that $e(S \cap L) = \binom{r}{2} - 2$, that is, S contains all except one of the edges e'' in L. Note that S necessarily contains all $r-2$ of the vertices in L that are not in e'. Therefore, by induction,

$$\mathscr{E}(S) \geqslant \mathscr{E}(S') + \lambda(r-2) - e(S \cap L) = \mathscr{E}(S') \geqslant 0.$$

Moreover, if equality holds then necessarily $\mathscr{E}(S') = 0$ and S' contains all vertices in e' that are not in e, as then $\sigma(S) = \sigma(S') + (r-2)$. In this case, $\bar{S}'$ is a subgraph of T' on the same vertex set as S'. It follows that $\mathscr{E}(\bar{S}') < \mathscr{E}(S') = 0$, which in turn implies that $\bar{S}' = T'$. Hence S contains all edges in T, except the branch $B = e''$ of T.

Intermediate cases. Finally, suppose that $0 < e(S \cap L) < \binom{r}{2} - 2$ and that S contains x many of the vertices in L that are not in e'. In this case,

$$\mathscr{E}(S) \geqslant \mathscr{E}(S') + \lambda x - e(S \cap L) > \mathscr{E}(S') \geqslant 0,$$

where the strict inequality follows by Lemma 3.20.

3.5.4. *Deficiency: Proof of Lemma 3.16(2).* Once again, the proof is by induction on k, in the same setting as in the proof of part (1) of the Lemma. The base case $k = 0$ is trivial, as then $S = \emptyset$, and so $P = \emptyset$ as well, in which case $\mathscr{E} = 0$ and $\mathscr{D} = 0$.

Recall that in our proof of Lemma 3.16(1) above, we showed that $\mathscr{E}(S)$ is greater than or equal to either $\mathscr{E}(S')$ or $\mathscr{E}(\bar{S}')$, depending on the case.

In the current proof, we will show that if there is no increment in $\mathscr{E}$, then there is no increment in $\mathscr{D}$. In addition, we observe that if $\mathscr{E}$ does increase, the its increment is bounded away from 0 by Lemma 3.20. This will complete the proof, noting that increments of $\mathscr{D}$ are all at most r.

Empty case. If S contains no edges in L and no vertices in L, except possibly those in e', then we have that $S = S' \subsetneq T'$, and so $\mathscr{E}(S) = \mathscr{E}(S')$ and $\mathscr{D}(S,T) = \mathscr{D}(S',T')$. Therefore, this case follows by the inductive hypothesis applied to S' and T'.

Complete case. Recall that, if $L \subset S$, then $\bar{S}' \subsetneq T'$ and $\mathscr{E}(S) = \mathscr{E}(\bar{S}')$. We claim that $\mathscr{D}(S,T) = \mathscr{D}(\bar{S}',T')$. Indeed, it is clear that the $r-2$ vertices of L that are not in e' belong to $F(S,T)$. Furthermore, for each $v \in e'$,

$$E_v(T) = (E_v(T') \setminus e') \cup E_v(L).$$

Therefore, for each such v, we have that $v \in P(S,T)$ if and only if $v \in P(\bar{S}',T')$, and so the result follows by the inductive hypothesis applied to $\bar{S}'$ and T'.

Nearly-complete case. Recall that when S contains all except one edge e'' of L, we showed that $\mathscr{E}(S) \geqslant \mathscr{E}(S')$ with equality if and only if S' contains all of the vertices in e' that are not in e. Suppose that equality holds, and consider the following two cases. First, if $\mathscr{E}(S') = \mathscr{E}(S) = 0$ then S contains all the edges of T except e'' whence $\mathscr{D}(S,T) \leqslant 2$ (if either of the vertices in e'' are in e then they are not accounted for in $\mathscr{D}(S,T)$ so there is inequality here). Second, if $\mathscr{E}(S') > 0$ then $\bar{S}'$ is also a proper subgraph of T' and $\sigma(S') = \sigma(\bar{S}')$ because S' contains all of the vertices in e' that are not in e. Therefore, $\mathscr{E}(S) = \mathscr{E}(S') = \mathscr{E}(\bar{S}') + 1$, and the claim follows by induction on $\bar{S}'$. On the other hand, recall that if $\mathscr{E}(S) > \mathscr{E}(S')$ then some $v \in e'$ (for which $v \notin e$) is not in S', in which case the increment $\mathscr{E}(S) - \mathscr{E}(S') \geqslant \lambda$.

Intermediate cases. Recall that, when $0 < e(S \cap L) < \binom{r}{2} - 2$, we noted that Lemma 3.20 implies $\mathscr{E}(S) > \mathscr{E}(S')$. However, this lemma in fact gives $\mathscr{E}(S) - \mathscr{E}(S') \geqslant \xi$, for some $\xi > 0$ depending only on r.

3.5.5. *Extendability: Proof of Lemma 3.16(3).* Finally, we turn to the third part of the lemma, regarding vertex-minimal tree extensions of S.

As in the previous parts, the proof is by induction on k. However, the case analysis will be different.

The base case $k = 0$ is trivial, so we proceed with the inductive step.

Case 1. If $e(S\cap L)=0$, then the claim follows directly by induction.

Case 2. If $\bar{S}'=T'$ then $\mathscr{E}(S')=0$, and so

$$\mathscr{E}(S)=\lambda x-e(S\cap L),$$

where, as before, x is the number of vertices in $S\cap L$ that are not in e'.

Taking $T_*=T$, we find that $\mathscr{T}(S)\leqslant r-2-x$, since $S\subsetneq T$ and $\sigma(T)=\sigma(S)+(r-2-x)$.

There are two sub-cases to consider:

Case 2a. If $e(S\cap L)=\binom{r}{2}-2$ then $x=r-2$ and $\mathscr{T}(S)=\mathscr{E}(S)=0$.

Case 2b. Otherwise, $0<e(S\cap L)<\binom{r}{2}-2$, since $\bar{S}'=T'$ and $S\subsetneq T$. Hence, by Lemma 3.20, we have that $\mathscr{E}(S)\geqslant\xi$, and so $\mathscr{T}(S)\leqslant c\mathscr{E}(S)$.

Case 3. Suppose that $e(S\cap L)>0$ and $\bar{S}'\subsetneq T'$. Note that

$$e(S)=e(\bar{S}')+e(S\cap L)-1\,,$$

where the -1 accounts for the removal of e'.

Denote by $\bar{T}_*$ the TWG for e satisfying $\bar{S}'\subsetneq\bar{T}_*$ and

$$\sigma(\bar{T}_*)=\sigma(\bar{S}')+\mathscr{T}(\bar{S}').$$

By induction, $\mathscr{T}(\bar{S}')\leqslant c\mathscr{E}(\bar{S}')$.

There are three sub-cases to consider:

Case 3a. If both of the vertices in e' are in $S\cup e$, then it follows that $\sigma(S)=\sigma(\bar{S}')+x$. Therefore,

$$\mathscr{E}(S)=\mathscr{E}(\bar{S}')+\lambda x-e(S\cap L)+1.$$

By Lemma 3.20, we find that $\mathscr{E}(S)\geqslant\mathscr{E}(\bar{S}')$, where equality holds if and only if $L\subset S$, and otherwise in fact $\mathscr{E}(S)\geqslant\mathscr{E}(\bar{S}')+1$.

We construct T_* from $\bar{T}_*$ by removing e' and grafting L in its place. As such, $v(T_*)=v(\bar{T}_*)+r-2$, and so

$$\mathscr{T}(S)\leqslant\mathscr{T}(\bar{S}')+(r-2-x).$$

Consequently, if $x=r-2$ then $\mathscr{T}(S)\leqslant c\mathscr{E}(S)$, and otherwise, we have that $\mathscr{T}(S)\leqslant\mathscr{T}(\bar{S}')+(r-3)$ and $\mathscr{E}(\bar{S}')\leqslant\mathscr{E}(S)-1$, and so $\mathscr{T}(S)\leqslant c\mathscr{E}(S)$.

Case 3b. If neither of the vertices in e' are in $S\cup e$, then $x\geqslant 2$, and equality holds if and only if $S\cap L$ is an edge. Furthermore,

$$\sigma(S)-\sigma(\bar{S}')=x-2,$$

where the -2 accounts for the vertices in e'. Therefore,

$$\mathscr{E}(S)-\mathscr{E}(\bar{S}')=\lambda(x-2)-e(S\cap L)+1.$$

We construct T_* as follows:

Case 3bi. If $x>2$, then note that $F=S\cap L$ has $3\leqslant v(F)=x<r$ and $e(F)=e(S\cap L)$. Hence (see Definition 3.18) $\mathscr{E}(S)-\mathscr{E}(\bar{S}')\geqslant\xi$. Therefore,

in this case, we can obtain T_* from $\bar{T}_*$ simply by removing e' and grafting L in its place. Then $\mathscr{T}(S) \leqslant \mathscr{T}(\bar{S}') + r - 3$, and the claim follows.

Case 3bii. Suppose that $x = 2$, so that $S \cap L$ is a single edge f. Then $\mathscr{E}(S) = \mathscr{E}(\bar{S}')$. In this case, we obtain T_* from $\bar{T}_*$ by a more delicate construction. The main observation is that e' is an isolated edge in $\bar{S}'$, since neither of the vertices in e' are in $S \cup e$. Likewise, f is an isolated edge in S. Hence, to obtain T_* from $\bar{T}_*$, we replace the edge e' in $\bar{T}_*$ with the edge f. We observe that:

- $S \subsetneq T_*$ (since neither vertex in e' is in an edge of S);
- T_* is a TWG for e (since neither vertex in e' is in e); and
- $v(T_*) = v(\bar{T}_*)$.

By induction, we find that

$$\mathscr{T}(S) \leqslant \mathscr{T}(\bar{S}') \leqslant c\mathscr{E}(\bar{S}') = c\mathscr{E}(S),$$

as required.

Case 3c. Finally, suppose that one vertex u in e' is in $S \cup e$ but that the other vertex v in e' is not. This case is somewhat similar to the previous Case 3b. For completeness, let us give the details.

First, note that $x \geqslant 1$,

$$\sigma(S) - \sigma(\bar{S}') = x - 1$$

and

$$\mathscr{E}(S) - \mathscr{E}(\bar{S}') = \lambda(x-1) - e(S \cap L) + 1.$$

We construct T_* as follows:

Case 3ci. If $x > 1$, we obtain T_* from $\bar{T}_*$ by removing e' and grafting L in its place. By similar reasoning as in Case 3bi, we find that $\mathscr{T}(S) - \mathscr{T}(\bar{S}') \leqslant r - 3$ and $\mathscr{E}(S) - \mathscr{E}(\bar{S}') \geqslant \xi$, and the claim follows.

Case 3cii. Otherwise, if $x = 1$ then $S \cap L$ is a single edge e'' between u (the vertex in e' that is in $S \cup e$) and some vertex v'' in L that is not in e'. In this case, we construct T_* from $\bar{T}_*$ by replacing v with v''. Note that

- $S \subsetneq T_*$ (since v is not in any edge of S);
- T_* is a TWG for e (since v is not in e); and
- $v(T_*) = v(\bar{T}_*)$,

and by induction

$$\mathscr{T}(S) \leqslant \mathscr{T}(\bar{S}') \leqslant c\mathscr{E}(\bar{S}') = c\mathscr{E}(S).$$

This concludes the proof of all parts in Lemma 3.16.

3.6. **Enumeration of partial TWGs.** Recall (see Lemma 3.6) that k-TWGs have $s+2$ vertices and $\lambda s+1$ edges, where $s=(r-2)k$ is its number of vertices that are not in e. By Lemma 3.8 (and (3.5)) there are

$$s^{O(1)}\gamma^s s!$$

many k-TWGs.

The following lemmas establish similar results for the number of partial TWGs S and the number of ways to complete such an S to a TWG T, up to an error depending on its edge efficiency parameter $\mathscr{E}(S)$.

Lemma 3.21 (Counting partial TWGs)**.** *The number of partial TWGs S, for a given edge e, with*

- *σ many vertices not in e; and*
- *edge efficiency $\mathscr{E}$*

is at most

$$\sigma^{c\mathscr{E}}\gamma^{\sigma}\sigma!,$$

for some constant c depending only on r.

Proof. By Lemma 3.16(3), every such partial TWG S is contained in some TWG $T_*=T_*(S)$ with $\sigma(T_*)\leqslant\sigma+c\mathscr{E}$. In particular, every such S is contained in a k-TWG with $k=\lceil(\sigma+c\mathscr{E})/(r-2)\rceil$. Therefore, we can construct such an S by first choosing such a k-TWG T and then selecting $\lambda\sigma-\mathscr{E}$ of its edges. By Lemma 3.8, the number of possible choices is at most

$$\gamma^{(r-2)k}[(r-2)k]!\binom{\lambda(r-2)k+1}{\lambda\sigma-\mathscr{E}}\leqslant\sigma^{c\mathscr{E}}\gamma^{\sigma}\sigma!,$$

as claimed. ■

Lemma 3.22 (Counting extensions of partial TWGs)**.** *Suppose that S is a partial TWG for a given edge e, with*

- *$\sigma=\sigma(S)$ many vertices not in e; and*
- *edge efficiency $\mathscr{E}=\mathscr{E}(S)$.*

Suppose also that V is a vertex set of size $(r-2)k+2$ that contains

- *all vertices in S;*
- *the edge e; and*
- *an additional $\sigma'=(r-2)k-\sigma$ other vertices.*

Then the number of k-TWGs T for e on V such that $S\subsetneq T$ is at most

$$k^{c(\mathscr{E}+1)}\gamma^{\sigma'}\sigma'!,$$

for some constants c depending only on r.

Proof. For a fixed S, there is a bijective correspondence between a k-TWG T for e containing S and the complement $\tilde{S}$ induced by all the edges in T that are not in S. Since $F(\tilde{S},T) = D(S,T)$ and $P(\tilde{S},T) = P(S,T)$, it follows that

$$\sigma(\tilde{S}) = (r-2)k - \sigma(S) + \mathscr{D}(S,T).$$

Hence, by Lemma 3.16(2),

$$\sigma(\tilde{S}) + \sigma(S) \leqslant (r-2)k + c\mathscr{E}(S) + 2.$$

Therefore, since $e(S) + e(\tilde{S}) = \lambda(r-2)k + 1$, we find that

$$\begin{aligned}\mathscr{E}(S) + \mathscr{E}(\tilde{S}) &= \lambda(\sigma(S) + \sigma(\tilde{S})) - (e(S) + e(\tilde{S}))\\ &\leqslant \lambda((r-2)k + c\mathscr{E}(S) + 2) - (\lambda(r-2)k + 1)\\ &\leqslant c(\mathscr{E}(S) + 1).\end{aligned}$$

Altogether, $\sigma(\tilde{S}) \leqslant \sigma' + c(\mathscr{E}(S)+1)$ and $\mathscr{E}(\tilde{S}) \leqslant c(\mathscr{E}(S)+1)$. Therefore, by applying Lemma 3.21, we find that there are at most

$$k^{c(\mathscr{E}+1)}\gamma^{\sigma'}\sigma'!$$

such $\tilde{S}$, which completes the proof. ■

3.7. **Sharp upper bound.** In this section, we prove the sharp upper bound Proposition 3.1.

To begin, let us return to the setting of Lemma 3.12, wherein we fix some $e \in E(K_n)$ and let $B_k^{(e)}$ be the number of balanced k-TWGs for e contained in $\mathscr{G}_{n,p}$. We now, more specifically, suppose that $k = \beta \log n$, for some $\beta > 0$ to be determined below (but will depend only on r).

Since $1 \ll \log n \ll \sqrt{n}$, Lemma 3.12 applies, and so

$$\mu = \mathbb{E}(B_k^{(e)}) \sim c_\beta (\log n)^{-c}(1+\varepsilon)^{(r-2)k} p, \tag{3.11}$$

as $n \to \infty$, for some constant $c_\beta > 0$ depending only on r and β (and, as always, c is some constant depending only on r).

Let us write $T_1 \sim T_2$ if T_1 and T_2 are two distinct balanced k-TWGs for e such that $e(T_1 \cap T_2) > 0$. We put

$$\Delta = \sum_{T_1 \sim T_2} \mathbb{P}(T_1, T_2 \subset \mathscr{G}_{n,p}) = \sum_{T_1 \sim T_2} p^{2\lambda(r-2)k+2-e(T_1\cap T_2)}.$$

We will use Janson's inequality, which we recall asserts that

$$\mathbb{P}(B_k^{(e)} = 0) \leqslant \exp\left(-\frac{\mu^2}{\mu+\Delta}\right).$$

In order to apply this inequality, we will first bound Δ/μ^2 from above. We will use the fact that

$$\frac{\mu^2}{\mu+\Delta} = \left(\frac{1}{\mu}+\frac{\Delta}{\mu^2}\right)^{-1}$$

and show that $1/\mu$ and Δ/μ^2 are both $\ll n^{-c}$, for some $c>0$.

Observe that

$$\begin{aligned}\Delta &= \sum_S \sum_{T_1\cap T_2=S} p^{2\lambda(r-2)k+2-e(S)}\\ &\leqslant \sum_S |\{T : S\subset T\}|^2 p^{2\lambda(r-2)k+2-e(S)},\end{aligned}$$

where the summation is over all non-empty partial TWGs S of balanced k-TWG for e (with labels in $[n]$) and, for each such S, $|\{T : S\subset T\}|$ is the set of all such T that contain S.

Then, by Lemmas 3.21 and 3.22, we find that

$$\begin{aligned}\Delta &\leqslant \sum_{\sigma,\mathscr{E}} \binom{n}{\sigma}\sigma^{c\mathscr{E}}\gamma^{\sigma}\sigma!\left(\binom{n}{\sigma'}k^{c(\mathscr{E}+1)}\gamma^{\sigma'}\sigma'!\right)^2 p^{2\lambda(\sigma+\sigma')+2-(\lambda\sigma-\mathscr{E})}\\ &\leqslant \sum_{\sigma,\mathscr{E}}(1+\varepsilon)^{\sigma+2\sigma'}p^{2+\mathscr{E}}k^{c(\mathscr{E}+1)},\end{aligned}$$

where the summation is over all σ and $\mathscr{E}$ for which there exists a partial TWG S of a balanced k-TWG for e with these specific parameters $\sigma(S)=\sigma$ and $\mathscr{E}(S)=\mathscr{E}$, and where, for convenience, we let $\sigma'=(r-2)k-\sigma$. Note that, by (3.11),

$$\mu \geqslant k^{-c}(1+\varepsilon)^{\sigma+\sigma'}p.$$

Therefore,

$$\frac{\Delta}{\mu^2} \leqslant k^c \sum_{\sigma,\mathscr{E}}(1+\varepsilon)^{-\sigma}(pk^c)^{\mathscr{E}}. \tag{3.12}$$

We claim that the summation in (3.12) is $o(n^{-c})$, for some c depending only on r. To prove this, we will take two cases, with respect to whether the partial TGW S is edge inefficient $\mathscr{E}>0$ or edge efficient $\mathscr{E}=0$. In the latter case, the assumption that T is balanced will finally come into play.

Case 1. There are $O(k^2)$ summands with $\mathscr{E}>0$. Note that, since $\mathscr{E}(S)=\lambda\sigma(S)-e(S)$, it follows that $\mathscr{E}\geqslant 1/(r-2)$. Also note that, trivially, $(1+\varepsilon)^{-\sigma}\leqslant 1$. As a consequence, recalling that $k=O(\log n)$, we deduce that the total contribution of these summands to (3.12) in this case is at most $n^{-(1+o(1))c}$, where $c=1/((r-2)\lambda)$, which only depends on r, as required.

Case 2. If $\mathscr{E}(S)=0$ for some partial TWG S of a TWG T then, by Lemma 3.16(1), S is obtained from T by removing a single branch B. Since we may assume that $S\neq\emptyset$, it follows that $B\neq S$. In other words, B is a branch of one of the $\binom{r}{2}-1$ primal branches of T. Therefore, S contains all vertices in the root $H^{(e)}$ of T and $\binom{r}{2}-2$ of the primal branches of T in full. Since T is balanced, each primal branch contains $(r-2)q$ vertices outside the root, were q is as in (3.6). Hence,

$$\sigma(S) \geqslant (r-2)+(r-2)q\left(\binom{r}{2}-2\right) \geqslant ck.$$

As such, the $O(k)$ summands in (3.12) in which $\mathscr{E}=0$ also have that $\sigma \geqslant ck$. Consequently, their total contribution is at most $k^c(1+\varepsilon)^{-ck}$. Recall that $k=\beta\log n$. Therefore, by taking β large, the total contribution is at most n^{-c}, for some $c>0$ depending only on r.

With the claim above now proved, we now have that

$$\Delta/\mu^2 \ll (k/n)^c,$$

for some $c>0$ depending only on r. Also note that by (3.11) and taking $\beta>0$ large, we may assume also that $\mu \gg n^c$.

Therefore, by Janson's inequality,

$$\mathbb{P}(B_k^{(e)}=0) \leqslant \exp(n^{-c}) \ll n^2.$$

Then, taking a union bound over all $e\in E(K_n)$, we find that, with high probability, $\mathscr{G}_{n,p}$ contains a balanced k-TWG for all such e. In particular, $\langle \mathscr{G}_{n,p}\rangle_{K_r}=K_n$, as claimed.

4. Towards the lower bound

In this section, we prepare for the proof of the upper bound in Theorem 1.1. Throughout, we let $H=K_r$, for some fixed $r\geqslant 5$. However, for convenience often write H to denote a copy of K_r.

4.1. **A sharp lower bounds for TWGs.**

Recall that in Section 3, we showed that if

$$p=\left(\frac{1+\varepsilon}{\gamma n}\right)^{1/\lambda},$$

for some $\varepsilon>0$, then, with high probability, all edges can be added to $\langle \mathscr{G}_{n,p}\rangle_{K_r}$ by a TWG of logarithmic size. The first observation we make towards a matching lower bound, is that TWGs are much less likely on the other side of this threshold.

Lemma 4.1 (TWG threshold)**.** *Fix some $e\in E(K_n)$. Let $\varepsilon>0$ and put*

$$p=\left(\frac{1-\varepsilon}{\gamma n}\right)^{1/\lambda}.$$

Then, with high probability, there is no TWG T for e in $\mathscr{G}_{n,p}$.

Proof. By Lemma 3.8, there are $O(\gamma^\sigma \sigma!)$ (labelled) TWGs for a given edge e of size $\sigma = (r-2)k$. Such TWGs have $e(T) = \lambda\sigma + 1$ many edges. Therefore, the probability that a given e has a TWG in $\mathscr{G}_{n,p}$ is bounded by

$$p \sum_{s=0}^{\infty} p^{\lambda s} \binom{n}{s} O(\gamma^s s!) \leqslant O(p) \sum_{s=0}^{\infty} (\gamma n p^\lambda)^s$$
$$= O(p) \sum_{s=0}^{\infty} (1-\varepsilon)^s = O(p/\varepsilon) \to 0,$$

and the result follows. ■

This result suggests that a sharp threshold for K_r-percolation occurs at $p_c \sim (\gamma n)^{-1/\lambda}$. However, TWGs are only the simplest possible type of witness graph (formally defined in the next section). To prove the sharp lower bound in Theorem 1.1 we will require a reasonably precise understanding of the structure of general witness graphs, and specifically, how far from tree-like they can be as a function of their edge density.

4.2. **Witness graph algorithm (WGA).** For each edge $e \in E(\langle G \rangle_H)$ eventually added to G by the H-dynamics, we assign a *witness graph (WG)* $W^{(e)} \subset G$ that bears witness to the addition of e, in the sense that $e \in E(\langle W^{(e)} \rangle_H)$. The graph $W^{(e)}$ is defined by the recursive procedure introduced in [9], which we call the *witness graph algorithm (WGA)*.

Definition 4.2 (Witness graphs)**.** Recall that E_t denotes the set of edges added to a graph G in the tth round of the H-dynamics. For each $e \in E(\langle W^{(e)} \rangle_H) = \bigcup_{t \geqslant 0} E_t$, we define $W^{(e)}$ as follows.

(1) If $e \in E_0$, we simply put $W^{(e)} = e$.
(2) Otherwise, if $e \in E_t$, for some $t \geqslant 1$, we select a copy $H^{(e)}$ of H, for which $E(H^{(e)} \setminus e) \subset \bigcup_{s<t} E_s$, and put

$$W^{(e)} = \bigcup_{f \in E(H^{(e)} \setminus e)} W^{(f)}. \tag{4.1}$$

Note that, in the above algorithm, there is potentially some flexibility in choosing the copy $H^{(e)}$ of H. However, any arbitrary choice will do, and will have no effect on our arguments.

4.2.1. *Aizenman–Lebowitz property.* As shown in [9], WGs have the following useful property, reminiscent of an analogous property for r-neighbor bootstrap percolation on $\mathbb{Z}^d$, first observed in that context by Aizenman and Lebowitz [1].

Definition 4.3 (Size of WG)**.** For convenience, we let $\sigma(W) = v(W) - 2$ denote the *size* of a WG. In other words, $\sigma(W)$ is the number of vertices in $W = W^{(e)}$ that are not in e.

Put

$$m(t) = \max_{e \in \bigcup_{s \leqslant t} E_s} \sigma(W^{(e)}).$$

Note that, if $E_0, E_1 \neq \emptyset$, then $m(0) = 0$ and $m(1) = v(H) - 2$. On the other hand, for $t > 1$, by the recursive structure of WGs,

$$m(t) \leqslant v(H) - 2 + (e(H) - 1)m(t-1) \leqslant e(H)m(t-1),$$

since $v(H) - 2 = m(1) \leqslant m(t-1)$. In other words, in each round of the H-dynamics, $m(t)$ expands by at most a factor of $e(H)$. As such, we have the following result; see Lemma 13 in [9] and Lemma 8 in [12].

Lemma 4.4 (A–L for WGs)**.** *Suppose that, for some edge e, its WG is of size $\sigma(W^{(e)}) \geqslant v(H) - 2$. Then, for all $\sigma' \in [v(H) - 2, \sigma(W^{(e)})]$, there is some edge f with WG of size $\sigma(W^{(f)}) \in [\sigma', e(H)\sigma']$.*

In order to identify the sub-critical region for K_r-percolation, we will use Lemma 4.4 as follows. If a given edge e is eventually added by the K_r-dynamics, then it has a witness graph W. Either W is of size at most $\beta \log n$, or else *some* edge f has a witness graph W' of size between $\beta \log n$ and $\binom{r}{2}\beta \log n$. Using these facts, our goal is to show that $\mathbb{P}(e \in E(\mathscr{G}_{n,p})) \to 0$.

4.3. **Red edge algorithm (REA).** We will use the following *red edge algorithm (REA),* as introduced in [9], to study non-tree WGs.

We assume that $e \notin E(G)$, so that $W^{(e)} \neq e$.

This algorithm describes the formation of $W = W^{(e)}$ for a given $e \in \langle G \rangle_H$. Essentially, the REA is obtained from the WGA (which recall describes the formation of *all* WGs) by suppressing all steps in the WGA that do not contribute to the formation of W. For technical convenience, we also "slow down" the dynamics, so that in each step a single copy of H is completed.

More formally, W can be described in terms of a series of copies of H,

$$H^{(e_1)}, H^{(e_2)}, \dots, H^{(e_m)},$$

such that (1) $e_m = e$ and (2) for all $i < j$, if $e_i \in E_t$ then $e_j \notin E_s$ for all $s < t$. In other words the e_i are weakly time ordered, with respect to the time of their addition to G. The edges e_i are distinct and not in G.

For ease of notation, let us write $H_i = H^{(e_i)}$. In each step $j \leqslant m$ of the REA, we add all edges in H_j that are not in $\bigcup_{i<j} H_i$. One such edge e_j is colored red. All other (if any) edges are in G and colored black. Note that the color of any other edges in H_j have already been assigned by the REA in

previous steps $i < j$. After all m steps of the REA have been completed, W is obtained as the graph induced by all the black edges.

It will sometimes be helpful to write H_i^* to emphasize that the REA defines an edge-colored version of H_i. In this way, W is simply obtained by removing all the red edges from $\bigcup_{i=1}^m H_i^*$.

4.3.1. *REA components.* The REA gives rise to a natural component structure, which evolves in time with the algorithm.

Recall that after the kth step of the REA, we have added k copies $H_1 \dots, H_k$ of H. Consider an auxiliary graph $\mathscr{A}_k$, with a vertex a_i for each H_i, with $i \leqslant k$. Two such vertices a_i and a_j, are neighbors if their corresponding H_i and H_j share at least one (red or black) edge.

Naturally, each connected component $\mathscr{C}$ of $\mathscr{A}_j$ corresponds to an edge-colored graph

$$C = \bigcup_{a_i \in \mathscr{C}} H_i^*,$$

which we refer to as an *component* existing after the jth step of the REA. Note that all such components are (red and black) edge disjoint (but can share vertices).

4.3.2. *REA step types.* We saw that a TWG T of size σ has $\lambda\sigma + 1$ edges. It turns out that when H is a clique (see [9, 12]) this is the minimal possible number of edges in a WG. As such we make the following definition.

Definition 4.5 (Excess edges)**.** We call

$$\chi(W) = e(W) - (\lambda\sigma(W) + 1) \geqslant 0.$$

the *excess number of (black) edges* in W.

In order to study the combinatorics of WGs with a given excess χ, it will be helpful to differentiate between various types of steps in the REA.

Definition 4.6 (Tree step (TS))**.** The ith step of the REA, when H_i is added, is called a *tree step (TS)* if it starts a new component or if all of the components being merged with H_i contain a single edge in H_i and are vertex disjoint outside of H_i.

Definition 4.7 (Internal red step (IntR))**.** The ith step of the REA, when H_i is added, is called an *internal red step (IntR)* if (i) H_i merges with a single existing component C, (ii) $V(H_i) \subset V(C)$, and (iii) only a red edge is added to C.

In other words, an IntR step can be the first step that occurs when running the K_r-dynamics internally on some component C.

Definition 4.8 (Costly steps)**.** The ith step of the REA is called a *costly step* if it is neither a TS nor an IntR. The associated *cost* κ_i to the step is $r = v(H)$ plus the number of vertices (if any) outside H_i that are in at least two of the components being merged with H_i. If the ith step is not costly (i.e., TS or IntR) then we put $\kappa_i = 0$. The *total cost* of W is

$$\kappa(W) = \sum_{i=1}^{m} \kappa_i.$$

The term *cost* is justified by the next result which follows directly by the proof of Lemma 12 in [12].

Lemma 4.9 (χ is $\Omega(\kappa)$)**.** *Let $H = K_r$, with $r \geqslant 5$. Then its excess number of (black) edges in W satisfies*

$$\chi(W) \geqslant \xi\,\kappa(W),$$

where $\xi = \xi(H) > 0$ is some constant, depending only on H.

In other words, each costly step is associated with an irreversible contribution to the excess χ.

We note that the above bound follows by the inductive proof of Lemma 12 in [12], which shows that this property holds for all components at any given time of the REA. (Naturally, we can extend the definitions of σ, χ and κ to components C.) Since there is a single component by the end of the process, the result follows.

Although there is a "price" to be paid for each costly step, they can also lead to the spread of (zero-cost) IntR steps. Indeed, the main challenge in analyzing general WGs is the delicate interplay between costly and IntR steps.

4.4. **Tree Components.** Our proof strategy for the lower bound hinges on the following "stability" property: a WG with a small excess is "close", in some combinatorial sense, to a TWG. A basic version of this property is proved in Section 5, and a more refined version in Section 6. For this purpose, we give the following key definitions:

Definition 4.10 (Tree components)**.** A component C which exists after some step of the REA, for a witness graph W, that was formed with only TSs is called a *tree component (TC)*.

The union of black and red edges in a TC is of the following form.

Definition 4.11 (K_r-tree)**.** We call a graph T a *K_r-tree* (or a *clique tree*) of *order* ϑ if it is the union of ϑ many copies $H_1, \ldots, H_\vartheta$ of K_r, and that, for every $1 < i \leqslant \vartheta$, H_i shares exactly one edge with $H_1 \cup \cdots \cup H_{i-1}$. We call an edge $e \in E(G)$ an *internal edge* of T if it is contained in at least two of the copies of K_r.

Remark 4.12. In Section 3, we used the letter T to refer to TWGs in the K_r dynamics. From now on, we use T to refer to K_r-trees. These structures are clearly related. Specifically, the union of red and black edges in a TWG is a K_r-tree. In general, every TC is a K_r-tree whose edges are colored red or black. Conversely, every K_r-tree can be attained as a TC via some edge coloring.

The following claim regarding the structure of K_r-trees will be useful throughout the remainder of this article.

Claim 4.13. *Suppose that a K_r-tree T is composed of ϑ copies $H_1,...,H_\vartheta$ of K_r (as in Definition 4.11). Then, we have the following.*

(1) Every clique in T is contained in some H_i.
(2) If vertices $x \neq y$ have at least three common neighbors in T then they are neighbors.

Proof. We prove this by induction on the order ϑ of T. The base case $\vartheta = 1$ is trivial. We assume that the claim holds for all K_r-trees of order $\vartheta - 1$ and that T is a K_r-tree of order ϑ. Let $T' = \cup_{i<\vartheta} H_i$. We use the following observation: if a vertex belongs to $V_\vartheta = V(H_\vartheta) \setminus V(T')$ then all its neighbors are in H_ϑ.

To prove the first claim, suppose that some vertices $v_1,...,v_j$ form a clique in T. If they are all in T' then the claim follows by induction. Otherwise, one of these vertices is in V_ϑ, whence all of them are in H_ϑ, as claimed.

For the second claim, consider x, y and their common neighbors $v_1,...,v_j$, for some $j \geqslant 3$. If they all belong to T' then the claim follows by induction. Therefore, at least one them is in V_ϑ. If a common neighbor, say v_1, is in V_ϑ, then $x, y \in V(H_\vartheta)$ and are therefore neighbors. Otherwise, we may assume that $x \in V_\vartheta$, whence $v_1,...,v_j$ belong to $V(H_\vartheta)$. At most 2 of these common neighbors lie in the intersection of H_ϑ and T', and since $j \geqslant 3$ at least one of them is in V_ϑ, in which case, as we have already shown, it follows that x, y are neighbors. ■

Lemma 4.14. *If T is a K_r-tree then $\langle T \rangle_{K_r} = T$, i.e., T is K_r-stable.*

Proof. If H is a copy of K_r and $H \setminus e \subset E(T)$ then the endpoints of e have $r - 2 \geqslant 3$ common neighbors and $e \in E(T)$ by Claim 4.13(2). Therefore, T is K_r-stable, as claimed. ■

4.5. **Target edges.** Note that, in the REA, all of the red edges e_i are added along the way towards eventually adding $e = e_m$. As such, we make the following definition.

Definition 4.15 (Target edges)**.** After the kth step of the REA for W, we call a red edge e_i, for some $i \leqslant k$, a *target edge* if it has not yet been *reused*, meaning that $e_i \notin E(H_j)$, for all $i < j \leqslant k$. We call $e = e_m$ the *final target edge* of W.

Note that, after the kth step, the most recent red edge e_k is a target edge, but there can be more. However, by induction (recalling (4.1) above), after the last mth step of the REA, the final target edge $e = e_m$ is the only remaining target edge. This observation, although quite simple, will play a crucial role in our arguments. For instance, we use it now to show that TWGs are the only WGs with no excess $\chi = 0$. This is a first step towards our goal of showing that a WG with a small excess is similar to a TWG.

Lemma 4.16. *If C is a TC that exists after some step of the REA with a unique target edge e then C is a TWG for e. In particular, if every step in the REA of a witness graph W is a TS, then W is a TWG.*

Proof. We proceed by induction of the number k of TSs used to form C. If $k = 1$ the claim is trivial. Otherwise, C is formed by a copy $H^{(e)}$ of K_r that has a single edge in common with each of the components $C_1,...,C_h$ that are being merged. Observe that $C_1,...,C_h$ are TCs, since C is a TC. Additionally, the edge $e_i \in E(C_i) \cap E(H)$, for each $1 \leqslant i \leqslant h$, is the unique target edge of C_i, as otherwise C would have another target edge besides e. Therefore, each C_i is a TWG for e_i. By Definition 3.2, we find that C is a TWG for e, as required. ■

Corollary 4.17. *For every witness graph W we have that $\chi(W) = 0$ if and only if W is a TWG.*

Proof. The fact that $\chi(W) = 0$ if W is a TWG is clear. In the other direction, recall that if $\chi = 0$ then all steps are either TS or IntR steps. However, in this case, if there was ever an IntR step during the REA, then the first such step occurs in a TC, in contradiction to Lemma 4.14. Therefore, W is a TC whose REA uses only tree steps, an so it is a TWG by Lemma 4.16. ■

5. Coarse lower bound

Since the full proof of the lower bound in Theorem 1.1 is somewhat involved, we will first give an argument that proves that $p_c = \Omega(n^{-1/\lambda})$. These arguments already contain some of the main ideas, and, together the upper bound in Section 3, answers Problem 3 in [9]. (We caution the reader that the proof of the sharp lower bound in Section 6 is not self-contained, but rather builds on the arguments in this section.)

Proposition 5.1 (Coarse lower bound)**.** *For $r \geqslant 5$, we have that*

$$p_c(n, K_r) = \Omega(n^{-1/\lambda}).$$

We note that Proposition 5.1 fails if $r = 4$, where p_c is of order $(n\log n)^{-1/\lambda}$ (as discussed in Section 1). We point out that the assumption that $r \geqslant 5$ is indeed used in Lemma 4.14 above and Claim 5.3 below, where both claims playing a key role in controlling the spread of IntR steps.

5.1. **Tree decomposition.** Consider the witness graph W for some edge e, as given by the REA. Recall (see Section 4.3.1) that after the jth step of the REA, we have an edge-disjoint union of components. Each such component C (of red and black edges) can be written as

$$C = \bigcup_{a_i \in \mathscr{C}} H_i^*,$$

where $\mathscr{C}$ is the connected component associated with C in the auxiliary graph $\mathscr{A}_j$. Recall that the H_i^* are the red and black edge-colored copies of K_r added in the REA.

By the following recursive procedure, we will decompose each such C into the following subgraphs: a *bad part* B (possibly empty) and some number of *tree parts* $T_1, \dots, T_k$ (possibly $k = 0$) . We will call this the *tree decomposition* of C. In the sequel we will show, by induction on the number of steps in the REA, that the decomposition of every component C satisfies the following properties:

(1) The underlying graph (red and black edges) of every tree part is a K_r-tree.
(2) If $B = \emptyset$ then $k = 1$ and $C = T_1$ is a TC.
(3) If $B \neq \emptyset$ then
   (a) the intersection graph $T_i \cap B$ of every tree part T_i and B consists of a single edge that is not a target edge of C (and is called the *base* of T_i); and
   (b) every two tree parts intersect in at most one vertex, and if they do, that vertex lies in B. In particular, the tree parts are pairwise edge disjoint.

Since these properties are proved inductively, we can (and will) use them to define the recursive procedure that constructs the tree decomposition.

For the base case of the recursive construction, suppose that C is a single copy of K_r with one red edge and $\binom{r}{2} - 1$ black edges. In this case, we simply put $B = \emptyset$ and $T_1 = C$.

Otherwise, let us suppose that C is formed by adding H_j^* and merging it with some previously defined components $C_1, \dots, C_h$. Each such C_i has a bad part B_i and tree parts $T_1^{(i)}, \dots, T_{k_i}^{(i)}$. The bad part B and tree parts $T_1, \dots, T_k$ of C are defined recursively, with respect to the following cases:

Tree step. If C is formed in a TS, we claim that there is at most one tree part in each C_i that contains an edge in H_j^*. Indeed, if B_i is empty then the claim is trivial by property (2). Otherwise, by property (3b) the tree parts of C_i are pairwise edge disjoint and a TS can intersect with C_i in at most one edge. Let T be the K_r-tree that is formed by H_j^* and the tree parts that intersect it (T is indeed a K_r-tree since $C_1,\dots,C_h$ are vertex-disjoint outside of H_j^*, since we are considering a TS), and let $T_1,\dots,T_s$ denote all the other tree parts in $C_1,\dots,C_h$. Consider two subcases:

- **Standard tree step.** If there is at most one non-empty bad part among $B_1,\dots,B_h$, then set $B := \bigcup_{i=1}^h B_i$ and the tree parts of C to be $T_1,\dots,T_s$ and the new tree part T.
- **Bad tree step.** Otherwise, there are at least two non-empty bad parts. In this case, we set $B := (\bigcup_{i=1}^h B_i) \cup T$ and the tree parts of C to be $T_1,\dots,T_s$.

We stress that in the case of a bad tree step, the K_r-tree T is *not* considered to be a tree part of C, but is rather used only as an *auxiliary tree* in the construction of its decomposition.

IntR step. If C is formed in an IntR step, then $h=1$ and all the edges in H_j^* except the new red edge e_j are already contained in C_1. In this case, we simply obtain B from B_1 by adding e_j, and the tree parts of C are those of C_1, unchanged.

Costly step. Finally, suppose that C is formed in a costly step. Let V_j be the vertex set containing $V(H_j)$ and all vertices outside of $V(H_j)$ that are in at least two of the C_i being merged with H_j^*. Recall that $\kappa_j = |V_j|$ is the cost of this step. We call V_j the set of vertices *involved* with this costly step.

In this case, we let B be the union of H_j^* with $\bigcup_{i=1}^h B_i$ and all tree parts T in some C_i such that $V(T)\setminus V(B_i)$ contains a vertex in V_j. The tree parts of C are all the tree parts T of the various C_i for which $V(T)\setminus V(B_i)$ is disjoint of V_j.

5.2. **Properties of the tree decomposition.** In this section, we prove the structural properties of the tree decomposition discussed in the previous section.

Lemma 5.2. *After every step of the REA for a witness graph W, the tree decomposition of any given existing component C satisfies the properties (1)–(3) listed in Section 5.1 above.*

Proof. The proof is by induction on the number of steps in the REA that are used to construct a component C.

Recall that, in the base case, C is a single copy of K_r, B is empty and $T_1 = C$. Therefore properties (1)–(3) follow directly.

Standard tree step. Suppose that C is formed in a standard TS. If all the B_i are empty, then by induction, every C_i is a TC, in which case C is a TC and properties (1)–(3) follow directly.

Otherwise, suppose that $B_1 \neq \emptyset$ and the other B_i are empty, whence $B = B_1$. In addition, $C_2,..,C_h$ are TCs using property (2) by induction, and are therefore contained in the new tree part T. Note that the intersection of H_j and C_1 is a single edge e, which is either in some tree part, T' say, of C_1, or else in B_1 but in no tree part of C_1. In consequence, the other tree parts $T_1,...,T_s$ are all the tree parts of C_1 besides T' (if it exists). In particular, since $T_1,...,T_s$ have not been effected by this step and $B = B_1$ has not changed, properties (1)–(3) hold for them by induction.

It remains to verify property (3) for T. Since $B = B_1 \subseteq C_1$ and the current step is a TS, the intersection of T and B is equal to the intersection of T' and B, if T' exists, and otherwise it is equal to the edge e. In former case, property (3a) holds for T by induction applied to T'. In the latter case, property (3a) holds, as the edge e has been reused to form T. Next, for property (3b), note that, similarly, the intersection of T and T_ℓ, for any $1 \leqslant \ell \leqslant s$, is equal to the intersection of T' and T_ℓ, if T' exists, and otherwise it is equal to the intersection of e and T_ℓ. In the former case, property (3b) holds by induction. In the latter case, by induction, T_ℓ intersects B in a single edge. This edge cannot be e, by our assumption that e is in no tree part of C_1. Therefore, e and T_ℓ intersect in at most one vertex, which belongs to B if it exists.

Bad tree step. Suppose that C is formed in a bad TS. Recall that, in this case, the auxiliary K_r-tree T is not a tree part of C, but is rather contained in its bad part B. As such, property (1) follows directly by induction since no new tree parts are created, and property (2) holds vacuously since $B \neq \emptyset$.

Also recall that, in this case, B is the union of $B_1,...,B_h$ and T. Let $1 \leqslant \ell \leqslant s$, and suppose that T_ℓ was a tree part of C_i, for some $1 \leqslant i \leqslant h$. Let X denote the intersection of T and C_i. As we saw in the previous case of a standard tree step, X is either a tree part T' of C_i, or a single edge in B_i. The vertex intersection of T_ℓ with B is comprised of its vertex intersection with $B_i \cup X$, since the current step is a tree step, and so no other vertex of B belongs to C_i. Furthermore, we claim that the intersection of T_ℓ with $B_i \cup X$ is equal to its intersection with B_i. Indeed, this is clear if X is an edge of B_i. Otherwise, if X is the tree part T', we derive the claim using property (3b) by induction, which asserts that the intersection of T_ℓ and T' is contained in B_i. Using this claim, property (3a) for T_ℓ follows by induction.

For property (3b), let $1 \leqslant l \neq l' \leqslant s$ and $1 \leqslant i, i' \leqslant h$ be given, such that T_ℓ is a tree part of C_i and $T_{\ell'}$ a tree part of $C_{i'}$. If $i = i'$, then (3b) follows directly by induction. Otherwise, since the current step is a TS, T_ℓ and $T_{\ell'}$ can intersect in at most one vertex from H_j which (if it exists) must belong to B, since $H_j \subset B$.

IntR step. Next, suppose that C is formed in an IntR step. Recall that, in this case, $h = 1$ and $H_j \setminus e_j \subset C_1$. We obtain B from B_1 by adding the new red edge e_j to B_1, and all (if any) tree parts T of C_1 are unchanged.

Therefore, as in the previous case, properties (1) and (3b) hold by induction, and property (2) holds vacuously. Hence, it remains only to verify property (3a). To this end, we will use the following claim to deduce that, for all tree parts T, the intersection $T \cap B$ is the same as $T \cap B_1$.

Claim 5.3. *Fix $r \geqslant 5$. Suppose that G, G' are graphs that satisfy $|V(G) \cap V(G')| \leqslant 2$. Then, the vertex set of every copy H of K_r that is completed by the K_r-dynamics on $G \cup G'$ is either contained in $V(G)$ or in $V(G')$.*

Proof. Denote $A = V(G) \setminus V(G')$ and $B = V(G') \setminus V(G)$. Towards a contradiction, let H be the first copy of K_r, adding some edge e to $G \cup G'$, that violates the claim. Observe that $H \setminus e$ has no edges between A and B, since no such edge was previously added by the K_r-dynamics. By our assumption on H, both $|V(H) \cap A|$ and $|V(H) \cap B|$ are positive. But then, since $r \geqslant 5$, one of these sets has at least two vertices, and so H has at least two edges between A and B, which is a contradiction. ■

Let T be a tree part of C. Note that, to show that $T \cap B = T \cap B_1$, we need to show that, when the new edge e_j is added to B_1, the tree part T does not acquire a new vertex in B. By Lemma 4.14, T is K_r-stable, and so at least one vertex in H_j is outside of T. By induction, by property (3), the graphs T and $(C_1 \setminus T) \cup B_1$ intersect in a single edge, which is the edge in $T \cap B_1$. Therefore, combining Claim 5.3 and $V(H_j) \not\subset V(T)$, we deduce that $V(H_j)$ is disjoint from $V(T) \setminus V(B_1)$. In particular, it follows that $T \cap B = T \cap B_1$, as required.

Costly step. Finally, suppose that C is formed in a costly step. Property (1) holds by induction, and (2) holds vacuously. Let T be a tree part of some C_i such that $(V(T) \setminus V(B_i)) \cap V_j = \emptyset$. Recall that such a T remains a tree part of C after the costly step. We claim that $V(T) \cap V(B) = V(T) \cap V(B_i)$. Indeed, suppose, towards a contradiction, that there exists a vertex v in $V(T) \cap V(B)$ that is not in $V(B_i)$. Then, in particular, $v \notin V_j$. More specifically, this means that v is not in $V(H_j)$ and does not belong to any other component $C_{i'}$ with $i' \neq i$. Therefore, v must lie in a tree part T' of C_i that was added to B. However, by induction, the intersection $T \cap T'$ is contained in B_i, which establishes the claim. Consequently, property (3a) follows by induction.

For property (3b), let T, T' be tree parts of C that belonged to $C_i, C_{i'}$ respectively prior to the step. If $i = i'$, then property (3a) holds directly by induction. Otherwise, $V(T) \cap V(T') \subset B_i \cap B_{i'}$. Indeed, every vertex of T not in B_i is not in V_j, hence it does not appear in $C_{i'}$ (and similarly for T'). Moreover, by induction, T intersects B_i in an edge, and likewise T' intersects $B_{i'}$ in an edge. These edges cannot be identical, since T and T' lie in different components. Therefore, they intersect in at most one vertex, which must belong to B, as required. ■

5.3. **Complexity bounds.** Our goal is to use the tree decomposition to bound the complexity of a witness graph W in terms of its cost. We start with the following claim.

Lemma 5.4. *Let C be a component existing after some step of the REA for a witness graph W. Let $\beta(C)$ denote the number of bad tree steps that occur during in the formation of C, and let $\kappa(C)$ be the cost of C. Then, $\beta(C) = O(\kappa(C))$.*

Proof. Let us consider the m' steps of the REA that are involved in the formation of C. Let b_j, for $1 \leqslant j \leqslant m'$, denote that number of components (in the REA restricted to C) that have a non-empty bad part after the jth step. Clearly, if a bad tree step occurs at time j then $b_j \leqslant b_{j-1} - 1$. On the other hand, $b_j \leqslant b_{j-1} + 1$ and equality can occur only if step j is a costly step. Therefore, the number of bad tree steps is bounded by the number of costly steps, which is $O(\kappa(C))$. ∎

Next, to describe the complexity of components C, we make the following definitions.

Definition 5.5 (Maximal tree parts)**.** We say that a tree part is *maximal* if it is inclusion maximal during in REA. We let $\omega = \omega(W)$ denote the total number of maximal tree parts.

Let us emphasize here that an axillary tree T added to B in a bad tree step is not itself a maximal tree part of W. Recall that such a T is formed by merging H_j^* with various tree parts T_i (such that at least two of these belong to a component C_i with a $B_i \neq \emptyset$). In this case, it is the various T_i that are maximal tree parts, rather than T.

Definition 5.6 (Tree number)**.** The tree number $\tau(C)$ of a component C is defined recursively. If C is a TC then $\tau(C) = 0$. Otherwise, if C is formed by merging some $C_1, \dots, C_h$ with H_j^*, we let $\tau(C) = \sum_{i=1}^{h} \tau(C_i) + \tau_j$, where τ_j is the number of tree parts in the various C_i that are added to the bad part B of C after the merge. When such a tree part T is added to B, we say that it has become *compromised*.

We note that there are two ways that a tree part can become compromised: (a) in a costly step it can interact with a vertex involved with the costly step (i.e., contain such a vertex outside of its bad part), or (b) in a bad tree step, it can be merged together with the copy of K_r added in this step to form the axillary tree used in such a step.

The following lemma bounds the tree number of a component in terms of its cost.

Lemma 5.7 (Complexity bound)**.** *Let C be a component that exists after some step of the REA for a witness graph W. Then,*

$$\tau(C) = O(\kappa(C)).$$

In particular, $\tau = O(\kappa)$, where $\tau = \tau(W)$ is the tree number of W and $\kappa = \kappa(W)$ is its total cost.

Proof. The lemma follows directly if C is a TWG, as then both sides are equal to 0. Otherwise, we prove that $\tau(C) = O(\kappa(C) + \beta(C))$, which suffices by Lemma 5.4.

Suppose that C is formed by merging some $C_1, \ldots, C_h$ with H_j^*. Then, by the definition of the cost of C, we have that

$$\kappa(C) = \sum_{i=1}^{h} \kappa(C_i) + \kappa_j,$$

where $\kappa_j = |V_j|$ is the cost of the step (where $\kappa_j > 0$ if a costly step, and $\kappa_j = 0$ otherwise).

To complete the proof, we show that if $\tau_j > 0$, then it is bounded by some constant factor multiplied by κ_j plus the increment in $\beta(C)$. Observe that $\tau_j > 0$ occurs only in two types of steps:

Bad tree step. In such a step, $\tau_j \leqslant h = O(1)$ and the number $\beta(C)$ of bad tree steps increases by 1.

Costly steps. In a costly step, for every C_i, each tree part of C_i that becomes compromised has a vertex in V_j that is not in the bad part B_i of C_i. In addition, each such compromised tree part contributes a different vertex to V_j by property (3b) of Lemma 5.2. Therefore, the contribution of C_i to τ_j is at most $|V_j|$. Since the number of C_i satisfies $h = O(1)$, it follows that $\tau_j = O(\kappa_j)$. ■

Finally, we deduce the following results, which asserts that the combinatorial complexity of a general WG (as measured in terms of the number and total size of its maximal tree parts) is bounded in terms of its cost κ.

Corollary 5.8. *If W is a TWG then $\omega(W) = 1$ and $\kappa(W) = 0$. Otherwise, more generally, its number of maximal tree parts satisfies $\omega(W) \leqslant \tau(W) + 1 = O(\kappa)$.*

Proof. If W is not a TWG then $\tau(W)$ is the number of maximal tree parts that ever become compromised (and added to a bad part) during the REA. Clearly, the final target edge of W can be a target edge of at most one maximal tree part.

We claim that every other maximal tree part T must be compromised during the REA. Towards a contradiction, suppose that some such T is never compromised. Suppose that after the ith step of the REA the tree part T is fully formed. At this time, it has at least one target edge, none of which are the final target edge of W.

Consider the jth step, for some $j > i$, that one of these target edges is reused. Such a time exists, since eventually all such edges will be reused. Note that, in the previous $(j-1)$th step, T is the tree part of some component, with a bad part B. By property (3a) in Lemma 5.2, each target edge of T has at least one vertex outside of B.

To complete the proof, we take cases with respect to the step type of the jth step:

Standard tree step. If the jth step is a standard tree step, then this would further expand T, and contradict its maximality.

Bad tree step. If the jth step is a bad tree step, then T would be part of an auxiliary K_r-tree that is added to some bad part, and so then T would, in fact, become compromised.

IntR step. In the proof of Lemma 5.2, we showed that, in an IntR step, all of the vertices in the copy of K_r that is completed in such a step belong to the bad part of the component that is being augmented. Therefore the jth step cannot be an IntR step since, as already mentioned, each target edge of T has at least one vertex outside of B.

Costly step. If the jth step is a costly step then, since each target edge of T has at least one vertex outside of B, it follows that at least one of the vertices in the target edge of T that is reused in this step belongs to V_j, and so T does, in fact, become compromised. ■

Corollary 5.9. *Let W be a witness graph of order σ and cost κ, and let $T_1,\dots,T_\omega$ be its maximal tree parts. Then,*

$$\sum_{s=1}^{\omega} |V(T_i)| \leqslant \sigma + O(\kappa)\,. \tag{5.1}$$

Proof. For $1 \leqslant s \leqslant \omega$, we say that a vertex $v \in V(T_s)$ is a base vertex of T_s if there exists a time in the REA in which T_s is a tree part of some component and v belongs to the bad part of this component. By Lemma 5.2 and the fact that the bad part of a component can only grow during the REA we find that every maximal tree part T_s has either 0 or 2 base vertices. Therefore, by Corollary 5.8, the base vertices contribute $O(\kappa)$ to the left side of (5.1).

Next, suppose that the same vertex v appears in two maximal tree parts $T \neq T'$, but is not a base vertex in either of them. We claim that there is a step in the REA in which distinct components $C \supset T$ and $C' \supset T'$ are merged, and one of the following holds: either this step is a bad tree step and $v \in V(H_j)$, or this step is a costly step and $v \in V_j$.

Before proving the claim, we observe that it implies the corollary. Indeed, each bad tree step j can contribute at most $|V(H_j)| = O(1)$ to $\sum_s |V(T_s)| - \sigma$, and by Lemma 5.4 there are $O(\kappa)$ such steps. In addition, each costly step j can contribute at most $O(\kappa_j)$, since each vertex in V_j can be in $O(1)$ components being merged in this step. Therefore, the total contribution of the costly steps is bounded by $O(\kappa)$.

We proceed with proving the claim. Let t (resp. t') be the first step in the REA in which v belongs to a tree part T_1 (resp. T_1') that is contained in T (resp. T'). Note that $t \neq t'$, since they are the times of standard tree steps, in which distinct tree parts are created. Assume that $t < t'$, and let C_1, C_1' denote the components containing T_1, T_1', resp., at time t'. Note that T_1' is a tree part of C_1' that will grow into T' in the future of the REA.

We claim that $C_1 \neq C_1'$. Suppose, towards a contradiction, that $C_1 = C_1'$. Recall that v is not a base vertex of T', and thus it is also not a base vertex of T_1'. In particular, v is not in the bad part of $C_1' = C_1$. That is, by time t', T_1 has grown into a tree part T_2 of C_1, where $T_1 \subseteq T_2 \subseteq T$, that has not yet been compromised. In consequence, we have that T_2 and T_1' are distinct tree parts of the same component, and they share the vertex v outside of its bad part, contradicting property (3b) of Lemma 5.2.

Therefore, there exists a step j in the REA in which components C, C' containing C_1, C_1', resp., are merged. If this is a costly step, then $v \in V_j$ since it belongs to at least two of the merged components. Otherwise, this has to be a tree step. In such a case, C, C' are disjoint outside of H_j, whence $v \in V(H_j)$. We conclude the proof by showing this step cannot be a standard tree step. Indeed, if it were so, the step would merge subgraphs of T, T' into the same tree part, in contradiction to T, T' being two distinct maximal tree parts. ■

5.4. **Proof of the coarse lower bound.** We start this section by applying Corollaries 5.8 and 5.9 to obtain a bound on the number of witness graphs of a given size and cost. by Corollary 5.8, a general W with total cost $\kappa > 0$ has only $\omega = O(\kappa)$ many maximal tree parts. By Lemma 4.9, such a W has an excess $\chi = \Omega(\kappa)$ number of edges. Hence, we want to show that the number (compared with TWGs of the same size) of such witness graphs is offset by the extra cost p^χ to be paid for any such W.

Recall that $\kappa = 0$ if and only if W is a TWG. By Lemma 3.8, the number of (labelled) TWGs for a given edge e of size $\sigma = (r-2)k$ is at most $O(\gamma^\sigma \sigma!)$. For general WGs, on the other hand, we will prove the following upper bound.

Lemma 5.10. *The number of (labeled) witness graphs W for a given edge e with size σ and cost $\kappa > 0$ is at most*

$$A^\sigma \cdot \sigma! \cdot \sigma^{O(\kappa)}$$

for some constant $A > 0$.

This bound is far from optimal, but suffices for our current purposes. Indeed, in our application (in the proof of Proposition 5.1 below) the size of W is $\sigma = O(\log n)$ and its cost $\kappa = O(\log\log n)$ is relatively small. As we will see, in this range, the above bound is already more than enough to prove Proposition 5.1.

Proof. Observe that every (black) edge in a witness graph belongs to one of its maximal tree components, or to $E(H_j)$ for some bad tree step or costly step j. Therefore, to construct the witness graph, we choose (i) the number ω and sizes $t_1,...,t_\omega$ of its maximal tree parts, (ii) unlabeled tree components $T_1,...,T_\omega$ of the chosen sizes, and (iii) the number and red/black edge-colorings of bad tree steps and costly steps.

Then, we consider a vertex-disjoint union X of these unlabeled tree components and a (colored, unlabeled) copy of K_r for each bad tree step and costly step. We conclude the construction by (iv) choosing an onto function from the unlabeled vertices of X to the labels $\{1,...,\sigma+2\}$ by which we identify the vertices from X and create the witness graph.

We turn to bound the number of witness graphs that can be constructed. (i) By Corollaries 5.8 and 5.9, we have that $\omega = O(\kappa)$ and $t_1+\cdots+t_\omega = \sigma + O(\kappa)$, hence there are $\sigma^{O(\kappa)}$ choices for the number and sizes of the TCs. (ii) There are at most A^x many TCs with x unlabeled vertices, for some $A>0$. Indeed, this follows directly by the tree-like structure of TCs. Therefore, there are $A^{\sigma+O(\kappa)}$ choices for the tree components. (iii) By Lemma 5.4 and the fact that the number of costly steps is bounded by κ, we have that there are $2^{O(\kappa)}$ choices for the number and edge-colorings of the copies of K_r corresponding to bad tree steps and costly steps. (iv) Finally, by the above considerations we also have that the total number of vertices in the disjoint union X is $\sigma+O(\kappa)$. Hence, the number of ways to label the vertices of X using all the labels from the set $\{1,...,\sigma+2\}$ is $\sigma!\cdot\sigma^{O(\kappa)}$, and this concludes the proof. ■

We will use the following technical result to deal with very costly WGs.

Lemma 5.11 (No very costly WGs)**.** *Fix some $e\in E(K_n)$. Let $a,b>0$ and put $p=(a/n)^{1/\lambda}$. Then, for some $c>0$, with high probability, there is no witness graph for e in $\mathscr{G}_{n,p}$ of size $\sigma \leqslant b\log n$ and excess $\chi \geqslant c\log\log n$.*

Proof. This follows by a simple union bound.

In fact, we will show that, for some sufficiently large $c>0$, with high probability, there are no subgraphs of $\mathscr{G}_{n,p}$ with $\sigma+2$ vertices and at least $\lambda\sigma+1+c\log\log n$ edges, within this range of $\sigma\leqslant b\log n$. Indeed, the expected number of such subgraphs is at most

$$\binom{n}{\sigma+2}\binom{\sigma^2/2}{\sigma+1+c\log\log n}p^{\lambda\sigma+1+c\log\log n}.$$

It can be seen (we omit the details) that this at most

$$n^{-(1+o(1))(c/\lambda-\lambda b)\log\log n},$$

and the result follows, taking c large and summing over the $O(\log n)$ possibilities for σ. ■

Finally, we are ready to prove the coarse lower bound $p_c = \Omega(n^{-1/\lambda})$.

Proof of Proposition 5.1. Let $p = (\varepsilon/n)^{1/\lambda}$, for some small $\varepsilon > 0$, to be determined below. Fix some edge $e \in E(K_n)$. We will show that $e \notin E(\langle \mathscr{G}_{n,p}\rangle_{K_r})$ with high probability.

First, we note that by Lemma 4.1, and taking $\varepsilon > 0$ sufficiently small, we may assume that if e has a WG in $\mathscr{G}_{n,p}$ then its cost $\kappa > 0$. Indeed, the only WGs with $\kappa = 0$ are TWGs by Corollary 4.17, and by Lemma 4.1 there are with high probability no TWGs for e in $\mathscr{G}_{n,p}$, provided that $\varepsilon < 1/\gamma$.

By the Aizenman–Lebowitz property, Lemma 4.4, it suffices to show that, for some $\beta > 0$, with high probability there are no WGs

- W for e with $\kappa(W) \geqslant 1$ and $\sigma(W)/\log n \leqslant \beta$; or
- W', for any edge f, with $\sigma(W')/\log n \in [\beta, \binom{r}{2}\beta]$.

Furthermore, note that, by Lemmas 4.9 and 5.11, we may assume that the cost κ of any such potential W or W' is $O(\log\log n)$.

To rule out the first event, we apply Lemmas 4.9 and 5.10. The expected number of such W is at most

$$\sum_{\sigma=0}^{\beta\log n}\sum_{\kappa=1}^{O(\log\log n)}\binom{n}{\sigma}A^{\sigma}\sigma!\sigma^{O(\kappa)}p^{\lambda\sigma+1+\xi\kappa} \leqslant (\log n)^{O(\log\log n)}\cdot p\cdot\sum_{\sigma=0}^{\beta\log n}(\varepsilon A)^{\sigma},$$

which tends to 0 as $n\to\infty$ provided $\varepsilon < 1/A$.

For the second event, the expected number of such W' is at most

$$\binom{n}{2}\sum_{\sigma=\beta\log n}^{\beta\binom{r}{2}\log n}\sum_{\kappa=0}^{O(\log\log n)}\binom{n}{\sigma}A^{\sigma}\sigma!\sigma^{O(\kappa)}p^{\lambda\sigma+1+\xi\kappa} \leqslant (\log n)^{O(\log\log n)}\cdot n^2p\cdot\sum_{\sigma\geqslant\beta\log n}(\varepsilon A)^{\sigma},$$

Note that here we include $\kappa = 0$ since W' is a WG for any edge f, hence we also need to consider the possibility it is a TWG. In this case, we use the fact that the number of TWGs of size σ is at most $A^\sigma \sigma!$. Clearly, for large enough $\beta > 0$, this expectation tends to 0 as $n \to \infty$, provided $\varepsilon < 1/A$. ■

6. Sharp lower bound

Finally, we now turn to the lower bound in Theorem 1.1.

Proposition 6.1 (Sharp lower bound)**.** *Fix $r \geqslant 5$ and let $\varepsilon > 0$. Put*

$$p = \left(\frac{1-\varepsilon}{\gamma n}\right)^{1/\lambda}$$

and fix $e \in E(K_n)$. Then, with high probability, $e \notin E(\langle \mathscr{G}_{n,p} \rangle_{K_r})$.

To prove this result, we build on the arguments of Section 5. A closer inspection of those arguments reveals why they fall short of establishing a sharp lower bound: The main issue is that, in bounding the number of WGs of a given size and cost, we relied on the estimate A^x for the number of TCs with x unlabeled vertices (see Lemma 5.10). However, $A > 0$ is a constant *strictly* larger than γ (the exponential growth rate of TWGs) since there are exponentially more TCs than TWGs.

To refine this approach, we must control the TCs that actually arise in the tree decomposition of a WG of *small* cost, and show that they are *nearly* TWGs. In this section, we achieve this by a detailed analysis of the spread of IntR steps within compromised tree components (i.e., those that have interacted with costly steps). Ultimately, we will prove that the number of target edges (a natural proxy for complexity) that a tree part of a WG can have is bounded, up to a constant factor, by the total cost of the WG.

More specifically, in this section, we will analyze the spread of new edges that occurs after a tree part T of a WG has interacted with other parts of the WG in costly (non-tree-like) ways. To be more concrete, suppose that T is a K_r-tree and G is some other graph. We want to bound the number of edges in $\langle G \cup T \rangle_{K_r}$ with both endpoints in T in terms of the size of the vertex set $S = V(G) \cap V(T)$. We think of T as a tree part and S as the vertices of T that are involved in costly steps in the REA. To bound this set of edges, we will make a comparison with a certain vertex bootstrap percolation process, which we call the $(r-2)_*$-BP process, defined in Section 6.1 below. In this context, S is the set of *seed vertices* that are initially infected. The *comparison lemma* (Lemma 6.4) in Section 6.2 shows that this process provides a useful comparison: all endpoints in edges that are added to T by the K_r-dynamics are eventually infected in $(r-2)_*$-BP. This initial comparison, however, is only one piece of the puzzle. In Section 6.3, we will prove an *expansion lemma* (Lemma 6.3) that limits the spread of infection under the $(r-2)_*$-BP dynamics in terms of the size of S. Finally, in Sections 6.4–6.6, we combine the comparison and expansion lemmas to deduce a *spread lemma* (Lemma

6.8), which is the key to proving a sharp lower bound on p_c. Using this result, we build upon the arguments in Section 5, obtaining very close control on the complexity of tree parts in a witness graph with a given total cost.

6.1. $(r-2)_*$**-BP on TCs.** Recall that the graph formed by the edges (both red and black) in a TC is a K_r-tree (see Definition 4.11.) We have found it difficult to directly analyze the spread of IntR steps throughout a TC after it interacts with costly steps. Instead, we control this spread indirectly, using the following modified version of the $(r-2)$-neighbor bootstrap percolation process on a K_r-tree.

In what follows, it will be helpful to think of T as some tree part in a WG, and the set $S \subset V(T)$ as the locations where it has interacted with costly steps. As we will see, the following vertex dynamics are chosen in such a way that if an edge is added to T by the K_r-dynamics, then both of its endpoints are eventually infected by the vertex dynamics. In our applications of this process, the infected vertices will represent locations where we lose control over the spread of edges by the K_r-dynamics, and so we will usually assume (the worst case scenario) that they induce a clique.

Definition 6.2 ($(r-2)_*$-BP)**.** Let T be a K_r-tree and $S \subset V(T)$ a set of *seeds*, which we consider to be initially *infected*. Then, in each step of the process, either

- *(usual step)* some vertex v with at least $r-2$ infected neighbors becomes infected; or else
- *(special step)* if no such v exists, but for some internal edge e, there are two cliques $H_i \neq H_j$ containing e such that H_i has $r-4$ infected vertices and H_j has 1 infected vertex, and all $r-3$ of these vertices are not in e, then we infect an arbitrarily chosen vertex $u \in e$; and otherwise,
- *(terminal step)* if neither type of step is possible, then the process terminates.

We call this process $(r-2)_*$*-BP* on T, and let $\langle S;T\rangle_*$ denote the set of eventually infected vertices in T started with the seed set S.

Remark 6.3. The reason why only one vertex $u \in e$ becomes infected by the special rule is for technical reasons that will become apparent below. Note that, after such a special step, the other vertex in e will become infected by the usual rule. Hence $(r-2)_*$-BP prioritizes usual steps, and uses special steps only when necessary.

As discussed in Section 2.4.2, the special step above is related to the presence of "bottlenecks" as new edges attempt to spread throughout a tree part T of a WG after it has interacted with costly steps. More specifically,

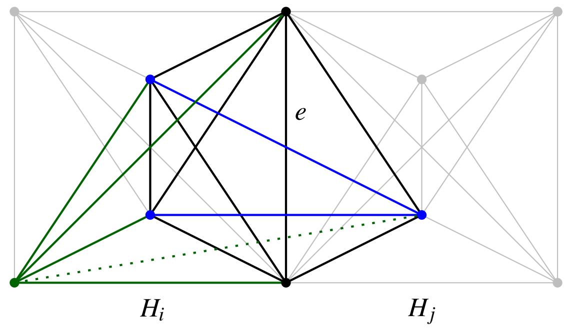

FIGURE 3. Illustration of the special step when $r = 6$. Two copies H_i, H_j of K_6 share an edge e. There are $r-3 = 3$ infected vertices (in blue), with $r-4 = 2$ in H_i and 1 in H_j. Assume that all other edges (in blue) between these vertices, that are not already in T, are eventually added by the K_r-dynamics. Then, these vertices together with those in e form a clique of size $r-1 = 5$ (the relevant edges of T are in black). The K_r-dynamics then add all other edges between H_i and the infected vertex in H_j; the addition of one such edge (dotted) is indicated with green edges. Likewise, once one vertex in e is infected by the special step, all other vertices in H_i (starting with the other vertex in e) get infected by usual steps. In this way, the special step is taken so as to maintain a useful comparison with the K_r-dynamics.

in the special step, if the $r-3$ infected vertices in $H_i \cup H_j$ form a clique (as discussed, in bounding the spread of edges, we will usually assume that the infected vertices induce a clique, since we have little control over the spread of edges between such vertices) then these vertices, together with the endpoints of the shared edge e, form a clique of size $r-1$. If this is the case, then the K_r-dynamics will add further (in fact, all) edges between H_i and the infected vertex in H_j. This is why we infect a vertex of e in the special step: since edges will begin to "flow" from one side of e to the other (i.e., between H_i and the infected vertex in H_j) we must (in order to maintain a useful comparison) allow the vertex dynamics to spread "through" this edge. Once one of the vertices in e is infected in the special step, all other vertices in H_i (including the other vertex in e) will be infected by usual steps, in line with the fact that the K_r-dynamics will add all edges between H_i and the infected vertex in H_j; see Figure 3. Furthermore, note that the configuration in the special step is *minimal*, in the sense that if there are fewer than $r-4$ infected vertices in H_i or none in H_j, then no more edges can be added between H_i and H_j by the K_r-dynamics, in which case e remains a "bottleneck" between H_i and H_j.

6.2. **Comparison lemma.** In order to study the interaction of costly steps with TCs, we will consider the following situation: Suppose we take the union of some graph G with a K_r-tree T and run the K_r-dynamics. Clearly, we cannot control the effect of the K_r-dynamics inside the graph G, since we did not assume anything on its structure. However, we will show that an edge e whose endpoints are in $V(T)$ can be infected only if both its endpoints are in the set of eventually infected vertices $\langle S;T\rangle_*$ in the $(r-2)_*$-BP process on T, if we take the set $S = V(G)\cap V(T)$ to be the seeds.

In our application of the following lemma, G will correspond to the graph that is comprised of the costly steps during the REA, which interact with a TC T. Roughly speaking, we will use it to show that IntR steps can only spread as far as $(r-2)_*$-BP.

Lemma 6.4 (Comparison lemma)**.** *Let $r \geqslant 5$, T be a K_r-tree, G a graph and $S = V(G)\cap V(T)$. Then,*

$$\langle G\cup T\rangle_{K_r} \subseteq Q\cup T\,,$$

where Q is a clique on the vertex set $V(G)\cup\langle S;T\rangle_*$.

We note that it is easy to see that $\langle G\cup T\rangle_{K_r} = Q\cup T$ if the graph G is K_r-percolating.

Proof. We prove the lemma using Claim 4.13 on the structure of K_r-trees. Denote $G_* = \langle G\cup T\rangle_{K_r}$ and $I_* = \langle S;T\rangle_*$. Suppose, towards a contradiction, that $\tilde e = \tilde u\tilde v \in E(G_*)\setminus E(Q\cup T)$ is the first such additional edge, with at least one vertex $\tilde v \in V(T)\setminus I_*$, that is added to G_* by the K_r-dynamics. Let $\tilde H$ be the copy of K_r that $\tilde e$ completes, such that $\tilde H\setminus \tilde e \subset G_*$. Denote $V(\tilde H) = \{\tilde u,\tilde v,x_1,\dots,x_{r-2}\}$. Note that since $\tilde e$ is the first additional edge added to G_* with an endpoint in $V(T)\setminus I_*$, and in addition, the vertex $\tilde v\notin I_*$, then the edges in $\tilde H$ containing $\tilde v$ are all in T. In particular, $V(\tilde H)\subset V(T)$, since $\tilde H$ is a clique.

First, note that if also $\tilde u\notin I_*$ then the $2(r-2)$ edges between $x_1,\dots,x_{r-2}$ and $\tilde u,\tilde v$ belong to T. Therefore, $\tilde u$ and $\tilde v$ have $r-2\geqslant 3$ common neighbors in T, and so, by Claim 4.13(2), they are neighbors in T, in contradiction to $\tilde e\notin T$.

Hence, we assume that $\tilde u\in V_*$. Suppose that, up to relabeling the x_i if necessary, that

$$\{x_1,\dots,x_{r-2}\}\setminus I_* = \{x_1,\dots,x_k\},$$

for some $0\leqslant k\leqslant r-2$. We note that all of the $r-2$ edges between $\tilde v$ and $x_1,\dots,x_{r-2}$ and all of the $k(r-1-k)$ edges between $x_1,\dots,x_k$ and $x_{k+1},\dots,x_{r-2},\tilde u$ are edges of T, since each such edge has at least one endpoint in $V(T)\setminus I_*$.

We consider the following cases:

Case 1. If $k \geqslant 3$ then $\tilde{u}$ and $\tilde{v}$ have at least three common neighbors in T, and by Claim 4.13(2) they are neighbors in T, in contradiction to $\tilde{e} \notin T$.

Case 2. If $k = 0$ then $\tilde{v}$ has $r-2$ neighbors in T that belong to I_*. This means that $\tilde{v}$ should have been infected in the $(r-2)_*$-BP process, and it contradicts the assumption that $\tilde{v} \notin I_*$.

Case 3. If $k = 1$ then x_1 has $r-2$ neighbors in T that belong to I_*: $x_2, \ldots, x_{r-2}$ and $\tilde{u}$. Similarly to Case 2, this contradicts $x_1 \notin I_*$.

Case 4. Finally, suppose that $k = 2$.

First, we claim that $\tilde{v}, x_1, \ldots, x_{r-2}$ induce a clique in T. Note that $\tilde{v}, x_1, x_2$ form a clique in T, and that each of these vertices are neighbors with all of $x_3, \ldots, x_{r-2}$. Moreover, by Claim 4.13(2), we see that $x_i x_j \in E(T)$, for all $3 \leqslant i < j \leqslant r-2$, since each such x_i, x_j have the three common neighbors $\tilde{v}, x_1, x_2$ in T. Therefore, by Claim 4.13(1), we find that $\tilde{v}, x_1, \ldots, x_{r-2}$ belong to some copy H_i of K_r in T.

Second, observe that $\tilde{u}, x_1, x_2$ also form a clique in T. Hence, by Claim 4.13(1), they belong to some copy H_j of K_r in T. Note that $H_i \neq H_j$, since $\tilde{e} \notin T$, and that H_i, H_j intersect in the edge $x_1 x_2$.

Finally, observe that $x_3, \ldots, x_{r-2} \in I_*$ are $r-4$ infected vertices in H_i, and $\tilde{u} \in I_*$ is 1 infected vertex in H_j. Therefore, either x_1 or x_2 would have been infected in $(r-2)_*$-BP in a special step (or in a usual step if one of them has more neighbors in I_*), in contradiction to $x_1, x_2 \notin I_*$. ■

Lemma 6.4 asserts that all the edges that are added to $G \cup T$ by the K_r-dynamics are in the clique Q of the vertices in $V(G) \cup \langle S;T \rangle_*$. However, in our application of this result, we actually need a slightly stronger statement, namely, that all the copies H of K_r that are completed during this run of the K_r-dynamics are fully contained (not only the edges that these H add to $G \cup T$) in Q.

Corollary 6.5. *Let $r \geqslant 5$, T be a K_r-tree, G a graph and $S = V(G) \cap V(T)$. Then, every copy $\tilde{H}$ of K_r that is completed by the K_r-dynamics started from $G \cup T$ satisfies $V(\tilde{H}) \subset V(G) \cup \langle S;T \rangle_*$.*

In other words, each vertex of such a $\tilde{H}$ is either in G or infected by the $(r-2)_*$-BP process on T started from S.

Proof. Denote $G_* = \langle G \cup T \rangle_{K_r}$ and $I_* = \langle S;T \rangle_*$. Suppose, in contradiction, that $\tilde{H}$ is a copy of K_r that is completed in the run of the K_r-dynamics on $G \cup T$ and $V(\tilde{H}) \cap V(T) \setminus I_* \neq \emptyset$. Let $\tilde{v}$ be a vertex in $V(\tilde{H}) \cap V(T) \setminus I_*$ and $\tilde{e} \in G_* \setminus (G \cup T)$ the edge that is added to complete $\tilde{H}$.

By Lemma 6.4, the endpoints of every edge in G_* that is not in T are in $V(G) \cup I_*$. Therefore, all the edges of $\tilde{H}$ that contain $\tilde{v}$ belong to T, and, in particular, $V(\tilde{H}) \subset V(T)$. Consequently, the endpoints of $\tilde{e}$ are in I_*, whence $\tilde{v} \notin \tilde{e}$.

We consider two cases:

First, if $1 \leqslant |V(\tilde{H}) \setminus I_*| \leqslant 2$ then $\tilde{v}$ has at least $r-2$ neighbors in T from I_*. This means that it should have been infected in the $(r-2)_*$-BP process, which would contradict the assumption that $\tilde{v} \notin I_*$.

Otherwise, suppose $|V(\tilde{H}) \setminus I_*| \geqslant 3$. The edges between $V(\tilde{H}) \setminus I_*$ and the endpoints of $\tilde{e}$ belong to T. Therefore, the endpoints of $\tilde{e}$ have at least 3 common neighbors in T. But then $\tilde{e} \in T$, by Claim 4.13(2), in contradiction to the assumption that $\tilde{e}$ is added by the K_r-dynamics on $T \cup G$. ■

6.3. **Expansion lemma.** As discussed above, we plan to use the comparison lemma (Lemma 6.4) to bound the spread of IntR steps after a tree part of a witness graph has interacted with costly steps. In doing so, we will need to gain some control on the size of $\langle S;T\rangle_*$, as a function of $|S|$.

Lemma 6.6 (Expansion lemma)**.** *Let T be a K_r-tree for some $r \geqslant 5$ and $S \subset V(T)$. Then*

$$|\langle S;G\rangle_*| = O(|S|).$$

In other words the $(r-2)_*$-BP process can expand at most linearly. In fact, in our application of this lemma, sub-exponential expansion would suffice, but we do not know a simpler proof of such a weaker result.

Before proceeding with the proof, which is somewhat subtle, let us give some intuition for why such a result should be true. The edge density of a K_r-tree T is approximately

$$\lambda' = \frac{r+1}{2} \approx \frac{1+\vartheta(\binom{r}{2}-1)}{2+\vartheta(r-2)}$$

if its order ϑ is large. If the infection spreads via $(r-2)_*$-BP mostly locally inside its copies of K_r, then the edge density of the graph $T_* \subseteq T$ induced by $\langle S;T\rangle_*$ will be lower (since K_r is strictly balanced). If, however, the infection spreads more globally, and the edge density of T_* is closer to λ', then many full copies of H will become infected. Intuitively, in this case, we expect that many times during the $(r-2)_*$-BP process, vertices will become infected by more than the minimal number $r-2$ of infected neighbors (note that, if a copy of K_r becomes fully infected, then its last vertex to become infected has at least $r-1$ infected neighbors at the time of its infection).

As in the work of Riedl [45], a key step in the following proof is to keep track of the number of infected neighbors that vertices have when they become infected. The added difficultly, in the current proof, is the presence of special steps, in which vertices can become infected with less than the usual threshold number of infection. A careful analysis is required to show that all such special steps are compensated for later on during the process.

Proof. Let S be a set of seeds for the $(r-2)_*$-BP process on a K_r-tree T. Let $H_1, \dots, H_\vartheta$ be the copies of K_r used in the construction of T (as in Definition 4.11). We also let

- $I_* = \langle S; T \rangle_*$ be the set of all eventually infected vertices;
- $I = I_* \setminus S$ be the set of all infected non-seed vertices; and
- E be the set of all edges with either both endpoints in I or one in S and the other in I.

Note that each vertex in I is incident to at least $r-3$ edges in E.

The proof follows three main parts.

Part 1. First, we relate $|E|$ and $|I|$ by carefully assigning edges to vertices at the time of their infection.

In each step of the infection process, when a vertex v is infected, there are at least $r-3$ edges in E between v and previously infected vertices. We call a step a *k-step* if there are exactly k such edges. Furthermore, we let t_k denote the total number of k-steps during the process.

Claim 6.7. *We have that*

$$|E| \geqslant \frac{r+1}{2}|I| + \frac{1}{8} \sum_{k \geqslant r-1} k t_k.$$

Since the inequality may not be intuitively clear, let us give some explanation before the proof. Recall that $(r+1)/2$ is the (asymptotic) edge density of a K_r-tree T. Consider the subgraph G_* induced by all eventually infected vertices in I_*. Intuitively, in order for the infection to spread, the seed set S will need to cover some of the less dense regions of T_*. As a result, the ratio of the number of edges in E by the number of eventually infected non-seeds in I will pick up a "boost" (the sum over $k \geqslant r-1$ in the claim), corresponding to times when vertices have strictly more than $r-2$ (at least $r-1$) infected neighbors (associated with denser regions of T_*).

Proof. Clearly,

$$|I| = \sum_{k \geqslant r-3} t_k$$

and

$$|E| = \sum_{k \geqslant r-3} k t_k.$$

Therefore,

$$\begin{aligned} |E| - \frac{r+1}{2}|I| &= \sum_{k \geqslant r-3} \left(k - \frac{r+1}{2}\right) t_k \\ &\geqslant \frac{r-7}{2} \cdot t_{r-3} + \sum_{k \geqslant r-1} \left(k - \frac{r+1}{2}\right) t_k \,. \end{aligned} \tag{6.1}$$

In this inequality we have used $r \geqslant 5$ to deduce that the (discarded) summand $(r-5)/2 \cdot t_{r-2}$, corresponding to $k = r-2$, is non-negative.

If $r \geqslant 7$ then the claim follows since the first term in (6.1) is non-negative and, in addition, $k-(r+1)/2 \geqslant k/8$ for every $k \geqslant r-1$ and $r \geqslant 7$.

For $r = 5,6$ we also need to bound t_{r-3} from above. Note that $(r-3)$-steps are special steps, in which one of the two uninfected vertices in some internal edge e becomes infected. This occurs due to the fact that there are exactly two $H_i \neq H_j$, both containing e, where H_i has exactly $r-4$ infected vertices and H_j has exactly one infected vertex (not in e, whose vertices are uninfected). Furthermore, recall that special steps are only taken when usual steps are not possible. Therefore, after this vertex in e becomes infected in a special step, all of the other vertices in H_i will become infected by usual steps. The step that infects the last vertex to become infected in H_i will be a k-step, for some $k \geqslant r-1$. Altogether, we find that

$$t_{r-3} \leqslant \sum_{k \geqslant r-1} t_k.$$

Consider the case $r = 6$. Combining this inequality with (6.1) yields

$$|E| - \frac{7}{2}|I| \geqslant \sum_{k \geqslant 5} (k-4)t_k\,,$$

and the claim follows since $k-4 \geqslant k/8$ for every $k \geqslant 5$.

In the case $r = 5$ we need a slightly better upper bound for t_2. In this case, since $r-4 = 1$, H_i and H_j play the same role and will both become entirely infected after the special step. Therefore, the two steps that infect the last vertex in each of them are a k-step for some $k \geqslant 4$. We deduce that

$$2t_2 \leqslant \sum_{k \geqslant 4} t_k.$$

Combining this inequality with (6.1) gives (in the case $r = 5$) that

$$|E| - 3|I| \geqslant \sum_{k \geqslant 4} (k-7/2)t_k\,,$$

and the claim follows since $k-7/2 \geqslant k/8$ for every $k \geqslant 4$. ■

Part 2. Next, we obtain some further inequalities by considering an exploration process of T, one H_i at a time.

To begin, we reveal one its edges. Then, in each subsequent step, we reveal an unexplored H_i which has exactly one edge in the subgraph of T revealed so far. Since T is a K_r-tree, all of T will eventually be revealed by this procedure. Also note that exactly $r-2$ vertices and $\binom{r}{2}-1$ edges in H_i are new, as the other 2 vertices are in a previously revealed edge.

Let n_j, for $0 \leqslant j \leqslant r-3$, be the number of steps in the exploration process when exactly j of the new vertices in the revealed H_i are in I_*. Quite crucially, we define n_{r-2} slightly differently, as the number of steps in which all $r-2$ new vertices (and hence all r vertices in H_i) are infected, but that not all vertices in H_i are seeds.

The lower bound

$$\sum_{j=1}^{r-2} jn_j \leqslant |I_*| \tag{6.2}$$

is clear. To obtain an upper bound, we note that the only other infected vertices that are not counted by this sum are either in the initial edge or in S. Hence

$$|I_*| \leqslant \sum_{j=1}^{r-2} jn_j + (2+|S|). \tag{6.3}$$

Next, we claim that

$$|E| \leqslant 1 + \sum_{j=1}^{r-2} \frac{j(j+3)}{2} n_j \tag{6.4}$$

Indeed, the 1 before the sum accounts for the possibility that the initial edge is in E. In addition, if j new vertices from I_* are exposed in a step then at most $\binom{j}{2} + 2j = j(j+3)/2$ new edges from E are exposed in that step.

Finally, we claim that

$$(r-2)n_{r-2} \leqslant \sum_{k \geqslant r-1} kt_k. \tag{6.5}$$

To see this, consider a step in which all $r-2$ of the new vertices in H_i are infected, but that not all vertices in H_i are seeds. Then, when the last vertex v in H_i is infected it has at least $r-1$ infected neighbors. If v is amongst the new vertices in this step then we obtain $r-1$ new edges in E; and otherwise, only $r-2$ such edges. The result follows, noting that each step in the $(r-2)_*$-BP process that such a vertex v becomes infected is a k-step, for some $k \geqslant r-1$. All of the above described edges, of which there are at least $(r-2)n_{r-2}$, contribute to the infection of these vertices.

Part 3. To conclude, we combine the above inequalities.

First, we multiply (6.2) by $(r+1)/2$, add it to (6.4), and rearrange, so as to obtain

$$|E| \leqslant 1 + \frac{r+1}{2}|I_*| - \sum_{j=1}^{r-3} \frac{j(r-2-j)}{2} n_j\,.$$

Together with Claim 6.7 and (6.5), this implies that

$$\frac{r+1}{2}|I| + \frac{r-2}{8} n_{r-2} \leqslant 1 + \frac{r+1}{2}|I_*| - \sum_{j=1}^{r-3} \frac{j(r-2-j)}{2} n_j,$$

and so

$$\sum_{j=1}^{r-3} \frac{j(r-2-j)}{2} n_j + \frac{r-2}{8} n_{r-2} \leqslant 1 + \frac{r+1}{2}|S|.$$

Multiplying this inequality by 8, and then adding $2+|S|$ to both sides, yields

$$\sum_{j=1}^{r-3} 4j(r-2-j)n_j + (r-2)n_{r-2} + (2+|S|) \leqslant 10 + (4r+5)|S|.$$

Finally, using (6.3) and the fact that $j \leqslant 4j(r-2-j)$, for every $1 \leqslant j \leqslant r-3$, we find that

$$|I_*| \leqslant 10 + (4r+5)|S| = O(|S|),$$

as required. ■

6.4. **Spread lemma.** Combining Corollary 6.5 and Lemma 6.6, we obtain the following result.

Lemma 6.8 (Spread lemma)**.** *Fix $r \geqslant 5$. Let T be a K_r-tree and G be a graph with $x = |V(T) \cap V(G)|$ many vertices in common with T. Let E_* be the set of edges in T that are contained in a copy of K_r completed by the K_r-dynamics started from $G \cup T$. Then we have that $|E_*| = O(x)$.*

In other words, if we start with T and add G, then at most $O(x)$ edges in T are ever used by the K_r-dynamics.

Proof. Denote $S = V(T) \cap V(G)$. By Corollary 6.5 the endpoints of the edges in E_* belong to $I_* := \langle S; T \rangle_*$. In addition, by Lemma 6.6, we have that $|I_*| = O(x)$. To conclude the proof, we observe that there are at most $O(|I_*|)$ edges in the subgraph of T induced by I_*. Indeed, when we expose the copies of K_r used in the construction of T, in every step there are $O(1)$ vertices and edges from that induced subgraph which are introduced. ■

6.5. **The structure of maximal tree parts.** Next, using Lemma 6.8, we obtain an upper bound on the complexity of the maximal tree parts of a witness graph. This observation will play a key role in our improvement on the combinatorial bound in Lemma 5.10, given in the next section.

Lemma 6.9. *Let W be a WG of cost $\kappa > 0$, and T be one of its maximal tree parts, as given by the tree decomposition of W (see Defintion 5.5). Then, T has at most $O(\kappa)$ target edges.*

Proof. Denote the size of W by σ. Recall that all the target edges of T, except perhaps one, must be used in the REA. Let G be the (typically non-disjoint) union of: (i) the other maximal tree parts $T_1, \ldots, T_{\omega-1}$ of W, and (ii) all the copies of K_r that are added in bad tree steps or costly steps.

Here we view G and T as edge-colored graphs that contain both black and red edges. To be precise, a tree part T_i contains all the edges in the tree steps that were used to create it, and they are assigned the color that was given to them by the REA. Similarly, each copy of K_r corresponding to a bad tree steps or costly steps contains all $\binom{r}{2}$ edges (some red and some black) according to the way the edges were colored in the corresponding step. Importantly, note that while all the black edges of W must belong to $G \cup T$, some of the red edges of W may be absent from $G \cup T$. This is the case for the red edges that are added in IntR steps, but are not used in subsequent tree steps or costly steps. However, all the red edges of W clearly belong to $\langle G \cup T \rangle_{K_r}$. Hence, every (red) target edge of T that gets used in the REA is either an edge of G or it belongs to a copy of K_r that is completed by the K_r-dynamics started from $G \cup T$.

Therefore, using Lemma 6.8, we complete the proof by showing that the intersection of T and G has $O(\kappa)$ vertices (and can therefore span only $O(\kappa)$ edges in the K_r-tree T). First, by Lemma 5.4 and the definition of the cost of a step, there are only $O(\kappa)$ vertices in the copies of K_r that are added in bad tree steps or costly steps. This accounts for the vertices in $T \cap G$ that participate in such steps. In addition, we derive

$$|V(T)| + \sum_{i=1}^{\omega-1} |V(T_i) \setminus V(T)| \geqslant \sigma - O(\kappa),$$

since every vertex that appears in some maximal tree part is accounted for at least once in the left side, and there are at most $O(\kappa)$ vertices in W that do not appear in any maximal tree part. By combining this inequality with Claim 5.9, we derive

$$\sum_{i=1}^{\omega-1} |V(T_i) \cap V(T)| \leqslant O(\kappa),$$

and the proof is concluded. ■

Claim 6.10. *The number of unlabeled tree components with x vertices and Δ target edges is at most $\gamma^x \cdot x^{O(\Delta)}$.*

Proof. Consider the target edges $e_1, \ldots, e_\Delta$ of a tree component T, listed in some arbitrary order. For each target edge e_i, let $T_i \subseteq T$ be a TWG for e_i, which can be found using the WGA. After T_i is found, we color all its edges in T black, so that all future TWGs T_j, with $j > i$, will share at most one edge with it. Consider the auxiliary graph, whose vertices are $T_1, \ldots, T_\Delta$, and such that two are adjacent if they share an edge. By induction, this auxiliary is a tree, and so

$$\sum_{i=1}^{\Delta} |V(T_i)| = x + 2(\Delta - 1).$$

Hence, to choose such a witness graphs there are $x^{O(\Delta)}$ choices for the sizes of the TWGs, at most $\gamma^{x+2(\Delta-1)}$ choices for the tree structure of the TWGs, by (3.5) above, and $x^{O(\Delta)}$ ways to glue together the TWGs into T. This concludes the proof. ■

6.6. **The WG bound, revisited.** With Lemma 6.9 in hand, we can significantly improve upon Lemma 5.10.

Lemma 6.11. *The number of (labelled) witness graphs W for a given edge e with size σ and cost $\kappa > 0$ is at most*

$$\gamma^{\sigma} \cdot \sigma! \cdot \sigma^{O(\kappa^2)}.$$

Proof. Recall that, in the proof of Lemma 5.10, we used the simple bound A^x for the number of TCs T with x unlabelled vertices. However, by Lemma 6.9, we now know that each of the $O(\kappa)$ many maximal tree parts T of W have $O(\kappa)$ many targets, and by Claim 6.10 there are at most $\gamma^x x^{O(\kappa)}$ such TCs.

We adapt the proof of Lemma 5.10 such that in item (ii) we choose for each T_i one of the (at most) $\gamma^{t_i} \cdot \sigma^{O(\kappa)}$ tree components with t_i vertices and $O(\kappa)$ target edges. Therefore, the total number of choices for item (ii) is

$$\prod_{i=1}^{\omega} \gamma^{t_i} \sigma^{O(\kappa)} = \gamma^{\sum_{i=1}^{\omega} t_i} \sigma^{O(\omega\kappa)} = \gamma^{\sigma+O(\kappa)} \sigma^{O(\kappa^2)},$$

where the last equality is by Corollaries 5.8 and 5.9. The rest of the argument is not changed and the result follows directly. ■

Proof of Proposition 6.1. We adapt the proof of Proposition 5.1 above. Once again, we fix some $e \in E(K_n)$ and show that $e \notin E(\langle \mathscr{G}_{n,p} \rangle_{K_r})$ with high probability.

Since

$$p = \left(\frac{1-\varepsilon}{\gamma n} \right)^{1/\lambda}$$

we have by Lemma 4.1 that with high probability e is not added by a TWG.

Hence, once again, by Lemma 4.4, it remains to rule out the existence of WGs

- W for e with $\kappa(W) \geqslant 1$ and $\sigma(W)/\log n \leqslant \beta \log n$; or
- W', for any edge f, with $\sigma(W')/\log n \in [\beta, \binom{r}{2}\beta]$,

where $\beta = \beta(\varepsilon) > 0$ will be determined below. As before, we can assume that any such W and W' has cost $O(\log\log n)$.

For the first event, we apply Lemmas 4.9 and 6.11. The expected number of such W is at most

$$\sum_{\sigma=0}^{\beta\log n}\sum_{\kappa=1}^{O(\log\log n)}\binom{n}{\sigma}\gamma^\sigma\sigma!\sigma^{O(\kappa^2)}p^{\lambda\sigma+1+\xi\kappa}$$
$$\leqslant (\log n)^{O((\log\log n)^2)}\cdot p\cdot\sum_{\sigma=0}^{\beta\log n}(1-\varepsilon)^\sigma,$$

which tends to 0 as $n\to\infty$.

For the second event, the expected number of such W' is at most

$$\binom{n}{2}\sum_{\sigma=\beta\log n}^{\beta\binom{r}{2}\log n}\sum_{\kappa=0}^{O(\log\log n)}\binom{n}{\sigma}\gamma^\sigma\sigma!\sigma^{O(\kappa^2)}p^{\lambda\sigma+1+\xi\kappa}$$
$$\leqslant (\log n)^{O((\log\log n)^2)}\cdot n^2p\cdot\sum_{\sigma\geqslant\beta\log n}(1-\varepsilon)^\sigma.$$

which tends to 0 as $n\to\infty$, provided $\beta>0$ is large enough. ■

7. Subcritical edge expansion

In this section we prove Theorem 1.2, which shows that, in the subcritical regime, the K_r-dynamics increases the edge density of $\mathscr{G}_{n,p}$ by a factor that converges to a deterministic *constant* $\rho=\rho(p)$. To accomplish this, we show that in the subcritical regime (i) almost every infected edge is infected by a unique "small" TWG, and (ii) the number of TWGs in $\mathscr{G}_{n,p}$ is, with high probability, close to $\rho\cdot p\binom{n}{2}$.

To compute the constant ρ we consider the generating function of the dth Fuss–Catalan numbers, given by

$$f(x)=\sum_{k=0}^{\infty}\mathrm{FC}_d(k)x^k.$$

Recall that we denote $\alpha_d=(d+1)^{d+1}/d^d$. By (3.2), f has radius of convergence $1/\alpha_d$, converging if and only if $|x|\leqslant 1/\alpha_d$. For such x, the recursion (3.3) implies that

$$xf(x)^{d+1}=f(x)-1. \tag{7.1}$$

It can be shown, for any fixed $0<x<1/\alpha_d$, that (7.1) has precisely two solutions, which are both greater than 1. Using the continuity of f and $f(0)=1$, one can deduce that $f(x)$ is the smaller of these two solutions.

In addition, $f(1/\alpha_d)=(d+1)/d$, since this is the unique solution for (7.1) when $x=1/\alpha_d$. Hence, we have that $1<f(x)<(d+1)/d$ for every $0<x<1/\alpha_d$ by monotonicity of f.

Proof of Theorem 1.2. Let

$$\rho=\sum_{k=0}^{\infty}\mathrm{FC}_{\binom{r}{2}-2}(k)(1/\bar{\alpha})^k.$$

By the discussion above, we find that ρ is the smallest root of the equation $\rho^{\binom{r}{2}-1}=\bar{\alpha}(\rho-1)$ and that it satisfies $1<\rho<\frac{\binom{r}{2}-1}{\binom{r}{2}-2}$.

Let $\varepsilon>0$, and choose a sufficiently large integer K so that

$$\rho_K:=\sum_{k\leqslant K}\mathrm{FC}_{\binom{r}{2}-2}(k)(1/\bar{\alpha})^k\in(\rho-\varepsilon,\rho). \tag{7.2}$$

We say that a TWG is *large* if its order is greater than K, and otherwise it is called *small*.

Denote by X_k the number of TWGs of order k that are contained in $\mathscr{G}_{n,p}$. Then,

$$\mathbb{E}(X_k)=\binom{n}{2}\binom{n-2}{(r-2)k}t(k)p^{\lambda(r-2)k+1}.$$

If k is fixed then

$$((r-2)k)!\binom{n-2}{(r-2)k}\sim n^{(r-2)k}.$$

Hence, using Lemma 3.8, it follows that

$$\mathbb{E}(X_k)\sim p\binom{n}{2}\cdot\mathrm{FC}_{\binom{r}{2}-2}(k)(1/\bar{\alpha})^k. \tag{7.3}$$

Let us turn to the variance of X_k. We claim that, for some constant $D>0$, we have that

$$\begin{aligned}\mathbb{E}(X_k^2)&\leqslant\mathbb{E}(X_k)+\mathbb{E}(X_k)^2+\sum_{v,e\geqslant 1}\binom{n}{2}^2 t(k)^2 n^{2(r-2)k-v}p^{2\lambda(r-2)k+2-e}\\&\leqslant\mathbb{E}(X_k)+\mathbb{E}(X_k)^2+D\left(p\binom{n}{2}\right)^2 n^{-v+e/\lambda}.\end{aligned}$$

In the first line, the summation is over all the integers $v,e\geqslant 1$ for which there exist k-TWGs T_1,T_2 that share e edges and v vertices. The first inequality holds since the first term accounts for pairs of identical TWGs, the second for pairs of edge-disjoint TWGs, and the third for pairs of TWGs that intersect non-trivailly. For the second inequality, we use the fact that $k,\bar{\gamma}$ are fixed and independent of n, hence $D>0$ is some constant.

Let T_1,T_2 be distinct k-TWGs for the edges e_1,e_2 respectively. Suppose that they share $e\geqslant 1$ edges and a set V of v vertices. We claim that $e<\lambda v$. Indeed, by Lemma 3.16(1) we have $e\leqslant\lambda|V\setminus e_1|\leqslant\lambda v$. Moreover, both inequalities cannot hold with equality, since if V is disjoint from e_1, then the intersection of T_1 and T_2 cannot be obtained from T_1 by removing one of its branches. Therefore, $\mathbb{E}(X_k^2) \sim \mathbb{E}(X_k)^2$ and so, by Chebyshev's inequality, $X_k \sim \mathbb{E}(X_k)$ in probability as $n \to \infty$. Since $K > 0$ is a constant, we find that

$$\sum_{k \leqslant K} X_k \sim \rho_K \cdot p \binom{n}{2}, \tag{7.4}$$

in probability as $n \to \infty$.

Next, to bound the contribution from large TWGs, we observe that the left side of (7.3) is bounded from above by the right side for every k (even if it grows with n). Therefore, the expected number of large TWGs in $\mathscr{G}_{n,p}$ is

$$\mathbb{E}\left(\sum_{k>K} X_k\right) < p\binom{n}{2}\varepsilon .$$

Therefore,

$$\mathbb{P}\left(\sum_{k>K} X_k > p\binom{n}{2}\sqrt{\varepsilon}\right) < \sqrt{\varepsilon}. \tag{7.5}$$

Let Y denote the number of pairs of small TWGs contained in $\mathscr{G}_{n,p}$ with the same target edge $f \in (K_n)$. Then,

$$\mathbb{E}(Y) \leqslant \binom{n}{2} \sum_{k_1,k_2 \leqslant K} \sum_{v,e} t(k_1)t(k_2) n^{(r-2)(k_1+k_2)-v} p^{\lambda(r-2)(k_1+k_2)+2-e},$$

where the second sum is over all the integers v, e for which there is a k_1-TWG and a k_2-TWG for the same edge f that have e edges and v vertices in common (besides the endpoints of f). Since $K, \bar{\gamma}$ are constants independent of n, we find that

$$\mathbb{E}(Y) \leqslant D \cdot \binom{n}{2} p^2 \cdot n^{-v+e/\lambda}.$$

for some constant D. By Lemma 3.16(1), $e \leqslant \lambda v$, whence $\mathbb{E}[Y] \leqslant D\binom{n}{2}p^2$, and

$$\mathbb{P}\left(Y > p\binom{n}{2}\varepsilon\right) \to 0, \tag{7.6}$$

as $n \to \infty$.

Let Z denote the number of edges $f \in K_n$ that can be added by the K_r-dynamics on $\mathscr{G}_{n,p}$ by a non-tree WG (with cost $\kappa \geqslant 1$). The proof of Proposition 6.1 shows that for some constant $\beta > 0$ the following holds:

(1) The expected number of edges that are added by a non-tree WG of size at most $\beta \log n$ is at most $\binom{n}{2} p^{1+o(1)}$.
(2) The probability that there exists an edge $f \in K_n$ that is added by a WG of size at least $\beta \log n$ tends to 0 as $n \to \infty$.

We conclude that
$$\mathbb{P}\left(Z > p\binom{n}{2}\varepsilon\right) \to 0\,, \tag{7.7}$$
as $n \to \infty$.

Finally, we observe that
$$\left| |E(\langle \mathscr{G}_{n,p}\rangle_{K_r})| - \sum_{k \leqslant K} X_k \right| \leqslant \sum_{k>K} X_k + Y + Z.$$
Indeed, on the one hand, an edge can belong to $\langle \mathscr{G}_{n,p}\rangle_{K_r}$ but not be a target of a small TWG only if it is the target of a large TWG or a non-tree WG. On the other hand, counting small TWGs may overestimate $\langle \mathscr{G}_{n,p}\rangle_{K_r}$, but only by at most the number Y of pairs of small TWGs sharing a target edge. Combining this observation with (7.2) and (7.4)–(7.7), we find that $|E(\langle \mathscr{G}_{n,p}\rangle_{K_r})|$ deviates by less than $(3\varepsilon + \sqrt{\varepsilon})p\binom{n}{2}$ from $\rho \cdot p\binom{n}{2}$ with probability at least $1 - \sqrt{\varepsilon} - o(1)$, and the proof is complete. ■

8. Final remarks

Many fascinating problems remain open.

8.1. **Near-critical structure.** Despite locating the first-order term of p_c, explaining “how" exactly percolation occurs remains open. Related to this, it would be interesting to investigate the critical window for K_r-percolation. For example, what is the second-order term in $p_c(n, K_r)$? In addition, simulations suggest that if the edges of K_n are added one-by-one in the so-called evolution of random graphs then there is, with high probability, a step m such that $\langle G(n,m)\rangle_{K_r}$ has density $O(n^{-1/\lambda})$ but $\langle G(n,m+1)\rangle_{K_r} = K_n$.

Let us also recall the discussion surrounding (2.1) above regarding a potential “critical droplet” of order L for K_r-percolation. It would appear that such a graph should have the peculiar property that it has no percolating proper subgraphs of size $k \gg 1$ (cf. the discussion in [9] at the top of p. 432 therein). We note, quite intriguingly, that if the K_r-dynamics were to somehow create a large enough clique of order L then, by the work of Janson, Łuczak, Turova and Vallier [35] (see Theorem 3.1 and compare (3.1) therein with (2.1) above), “nucleation” (that is, the $(r-2)$-neighbor dynamics) would then take over and lead to full percolation.

Finally, let us mention that another natural direction for future study is the speed (number of rounds) of the K_r-dynamics on $\mathscr{G}_{n,p}$. Such questions might be the most interesting when p is barely supercritical (but can be asked more generally), as the speed will be intimately related to the (still largely mysterious) way in which $\mathscr{G}_{n,p}$ percolates. We note that the speed of the r-neighbor dynamics on $\mathscr{G}_{n,p}$ is studied in [35].

8.2. **Slow convergence.** The simulations in Figure 1 above might give the impression that K_5-percolation spreads by nucleation, however, this is not be representative of the behavior as $n \to \infty$; let us explain. In these simulations, $n = 2000$ and $r = 5$, and so L in (2.1) is only about $L \approx 1.6$. In this figure we are, in fact, mostly seeing the nucleating 3-neighbor dynamics.

Bootstrap percolation is notoriously difficult to simulate. For instance, as discussed in [31, 34], conjectures in the physics literature for the constant in $p_c \sim (\pi^2/18)/\log n$ for the 2-neighbor dynamics on $[n]^2$ were far from the true value. That being said, simulations with only $n = 300$ (see, e.g., [33]) already show the correct qualitative super-critical behavior of the model: a rectangle grows by nucleation (its perimeter slowing expanding) until it reaches a critical size (a "critical droplet"), and then the process explodes.

In this respect, K_r-percolation, with $r \geqslant 5$, seems to be even more impervious to simulation: We expect nucleation (the $(r-2)$-neighbor dynamics) to take over once a clique of size order L is formed. However, it appears that n would need to be extremely large (perhaps well outside any computationally feasible range) in order for L to be large enough to provide some clues about how $\mathscr{G}_{n,p}$ manages to form this "critical droplet" of size L.

This is perhaps reminiscent of the *slow aging* phenomenon in statistical physics [32], where a system converges to its equilibrium very slowly, and at various intermediate time scales can exhibit different types of behavior. In our current context, K_r-percolation, with $r \geqslant 5$, and K_4-percolation appear to behave similarly (the $(r-2)$-neighbor dynamics dominate) for reasonably large values of n, but nonetheless, their behaviors diverge quite drastically in the large n limit.

8.3. **Strictly balanced and general graph templates.** Our work focuses on the case $H = K_r$, but many questions about H-percolation remain open for general graphs H. Most prominently, it would be interesting to determine $p_c(n,H)$ and whether the H-percolation threshold is sharp.

Cliques K_r are balanced, but only strictly balanced once $r \geqslant 5$, and this property plays a central role in our proof of Theorem 1.1. Hence, it seems plausible that $p_c = \Theta(n^{-1/\lambda})$ for *all* strictly-balanced graphs H. The upper bound $p_c = O(n^{-1/\lambda})$ for such H is proved in [12]. Moreover, it is possible that our proof can be adapted to all strictly-balanced graphs Indeed, our proof of the sharp upper bound in Proposition 3.1 does not use the fact that H is a clique in any crucial way (only that K_r, $r \geqslant 5$, is strictly balanced), apart from the fact that the exponential growth rate γ of TWGs has a convenient form in this case. Our proof of the sharp lower bound in Proposition 6, on the other hand, uses the fact that H is a clique in a number of places. However, we note that, in our proof of Proposition 6.1, there is an extra factor $p^{\xi\kappa}$ that is not fully used. In fact, we only used here that $\xi \geqslant 0$, although in fact

$\xi > 0$ by Lemma 4.9. Indeed, the only place in our proof that $\xi > 0$ plays a role is in ruling out very costly witness graphs. It could be that this extra factor $p^{\xi\kappa}$ will be helpful in analyzing general strictly balanced graphs H.

8.4. **Rigidity of witness graphs.** The connection between K_r-percolation and the algebraic notion of generic $(r-2)$-rigidity theory goes back to Kalai [36]. A similar connection between volume-rigidity and hypergraph bootstrap percolation was used in [42] to obtain a lower bound for the corresponding p_c. Hence, it seemed plausible to us, at the outset of this work, that rigidity theory would be helpful with studying K_r-percolation, but, in the end, we were not able to make progress with this approach. For instance, we are unable to prove the following stronger version of the spread lemma (Lemma 6.8): Fix $r \geqslant 5$. Let T be a K_r-tree, G be a graph with $|V(G) \cap V(T)| \leqslant x$, and C denote the closure in the generic $(r-2)$-rigidity matroid of $G \cup T$. Then, is it true that at most $O(x)$ vertices of T have a neighbor in C that was not already their neighbor in T? This is stronger than the spread lemma since $C \supseteq \langle G \cup T \rangle_{K_r}$, which follows from [36].

Acknowledgements

ZB was supported by National Research, Development and Innovation Office (NKFIH) grant KKP no. 138270. BK was partially supported by a Florence Nightingale Bicentennial Fellowship (Oxford Statistics) and a Senior Demyship (Magdalen College). GK was supported by the Royal Commission for the Exhibition of 1851. YP was partially supported by the Israel Science Foundation grant ISF-3464/24. Part of this project was completed during a visit to Magdalen College in July 2024. We thank the college for its hospitality and wonderful working environment.

References

1. M. Aizenman and J. L. Lebowitz, *Metastability effects in bootstrap percolation*, J. Phys. A **21** (1988), no. 19, 3801–3813.
2. N. Alon, *An extremal problem for sets with applications to graph theory*, J. Combin. Theory Ser. A **40** (1985), no. 1, 82–89.
3. O. Angel and B. Kolesnik, *Sharp thresholds for contagious sets in random graphs*, Ann. Appl. Probab. **28** (2018), no. 2, 1052–1098.
4. ______, *Large deviations for subcritical bootstrap percolation on the Erdős–Rényi graph*, J. Stat. Phys. **185** (2021), no. 2, Paper No. 8, 16.
5. A. C. Antonir, Y. Peled, A. Shapira, M. Tyomkyn, and M. Zhukovskii, *When does a tree activate the random graph?*, preprint available at arXiv:2507.05697.
6. M. Akhmejanova, I. Vorobyev, and M. Zhukovskii. *Weak saturation numbers of large complete bipartite graphs*, preprint available at arXiv:2508.19435.
7. P. Balister, B. Bollobás, R. Morris, and P. Smith, *Universality for monotone cellular automata*, J. Eur. Math. Soc. (2025), in press.

8. J. Balogh, B. Bollobás, H. Duminil-Copin, and R. Morris, *The sharp threshold for bootstrap percolation in all dimensions*, Trans. Amer. Math. Soc. **364** (2012), no. 5, 2667–2701.
9. J. Balogh, B. Bollobás, and R. Morris, *Graph bootstrap percolation*, Random Structures Algorithms **41** (2012), no. 4, 413–440.
10. J. Balogh, B. Bollobás, R. Morris, and O. Riordan, *Linear algebra and bootstrap percolation*, J. Combin. Theory Ser. A **119** (2012), no. 6, 1328–1335.
11. J. Balogh, G. Kronenberg, A. Pokrovskiy, and T. Szabó, *The maximum length of K_r-bootstrap percolation*, Proc. Amer. Math. Soc., to appear, preprint available at arXiv:2507.1907.04559.
12. Z. Bartha and B. Kolesnik, *Weakly saturated random graphs*, Random Structures Algorithms **65** (2024), no. 1, 131–148.
13. Z. Bartha, B. Kolesnik, and G. Kronenberg, *H-percolation with a random H*, Electron. Commun. Probab. **29** (2024), Paper No. 33, 5 pp.
14. E. Bayraktar and S. Chakraborty, *$K_{r,s}$ graph bootstrap percolation*, Electron. J. Combin. **29** (2022), no. 1, Paper No. 1.46, 17.
15. M. Bidgoli, A. Mohammadian, and B. Tayfeh-Rezaie, *On $K_{2,t}$-bootstrap percolation*, Graphs Combin. **37** (2021), no. 3, 731–741.
16. M. Bidgoli, A. Mohammadian, B. Tayfeh-Rezaie, and M. Zhukovskii, *Threshold for stability of weak saturation*, J. Graph Theory **106** (2024), no. 3, 474–495.
17. B. Bollobás, *Weakly k-saturated graphs*, Beiträge zur Graphentheorie (Kolloquium, Manebach, 1967), B. G. Teubner Verlagsgesellschaft, Leipzig, 1968, pp. 25–31.
18. B. Bollobás, M. Przykucki, O. Riordan, and J. Sahasrabudhe, *On the maximum running time in graph bootstrap percolation*, Electron. J. Combin. **24** (2017), no. 2, Paper No. 2.16, 20 pp.
19. R. Cerf and E. N. M. Cirillo, *Finite size scaling in three-dimensional bootstrap percolation*, Ann. Probab. **27** (1999), no. 4, 1837–1850.
20. J. Chalupa, P. L. Leath, and G. R. Reich, *Bootstrap percolation on a Bethe lattice*, J. Phys. C **21** (1979), L31–L35.
21. S. Diskin, I. Hoshen, D. Korándi, B. Sudakov, and M. Zhukovskii, *Saturation in random hypergraphs*, preprint available at arXiv:2405.03061.
22. P. Erdős and A. Rényi, *On random graphs. I*, Publ. Math. Debrecen **6** (1959), 290–297.
23. A. Espuny Díaz, B. Janzer, G. Kronenberg, and J. Lada, *Long running times for hypergraph bootstrap percolation*, European J. Combin. **115** (2024), Paper No. 103783, 18 pp.
24. R. J. Faudree, R. J. Gould, and M. S. Jacobson, *Weak saturation numbers for sparse graphs*, Discuss. Math., Graph Theory **33** (2013), no. 4, 677–693.
25. D. Fabian, P. Morris, and T. Szabó, *Slow graph bootstrap percolation II: Accelerating properties*, J. Combin. Theory Ser. B **172** (2025), 44–79.
26. U. Feige, M. Krivelevich, and D. Reichman, *Contagious sets in random graphs*, Ann. Appl. Probab. **27** (2017), no. 5, 2675–2697.
27. P. Frankl, *An extremal problem for two families of sets*, European J. Combin. **3** (1982), no. 2, 125–127.
28. M. Gardner, *Mathematical games: The fantastic combinations of John Conway's new solitaire game "life"*, Sci. Am. **223** (1970), no. 4, 120–123.
29. K. Gunderson, S. Koch, and M. Przykucki, *The time of graph bootstrap percolation*, Random Structures Algorithms **51** (2017), no. 1, 143–168.

30. I. Hartarsky and L. Lichev, *The maximal running time of hypergraph bootstrap percolation*, SIAM J. Discrete Math. **38** (2024), no. 2, 1462–1471.
31. I. Hartarsky and A. Teixeira, *Bootstrap percolation is local*, preprint available at arXiv:2404.07903.
32. M. Henkel and M. Pleimling, *Non-equilibrium phase transitions. Volume 2: Ageing and dynamical scaling far from equilibrium*, Theoretical and Mathematical Physics, Springer, Dordrecht, 2010.
33. A. E. Holroyd, *Bootstrap percolation*, https://personal.math.ubc.ca/~holroyd/boot/, webpage accessed on 2025-10-26.
34. ______, *Sharp metastability threshold for two-dimensional bootstrap percolation*, Probab. Theory Related Fields **125** (2003), no. 2, 195–224.
35. S. Janson, T. Łuczak, T. Turova, and T. Vallier, *Bootstrap percolation on the random graph* $G_{n,p}$, Ann. Appl. Probab. **22** (2012), no. 5, 1989–2047.
36. G. Kalai, *Weakly saturated graphs are rigid*, Convexity and graph theory (Jerusalem, 1981), North-Holland Math. Stud., vol. 87, North-Holland, Amsterdam, 1984, pp. 189–190.
37. ______, *Hyperconnectivity of graphs*, Graphs Combin. **1** (1985), no. 1, 65–79.
38. B. Kolesnik, *The sharp* K_4*-percolation threshold on the Erdős–Rényi random graph*, Electron. J. Probab. **27** (2022), Paper No. 13, 23 pp.
39. D. Korándi and B. Sudakov, *Saturation in random graphs*, Random Structures Algorithms **51** (2017), no. 1, 169–181.
40. G. Kronenberg, T. Martins, and N. Morrison, *Weak saturation numbers of complete bipartite graphs in the clique*, J. Combin. Theory Ser. A **178** (2021), Paper No. 105357, 15 pp.
41. L. Lovász, *Flats in matroids and geometric graphs*, Combinatorial surveys (Proc. Sixth British Combinatorial Conf., Royal Holloway Coll., Egham, 1977), Academic Press, London-New York, 1977, pp. 45–86.
42. E. Lubetzky and Y. Peled, *The threshold for stacked triangulations*, Int. Math. Res. Not. IMRN (2023), no. 19, 16296–16335.
43. R. Morris, *Monotone cellular automata*, Surveys in combinatorics 2017, London Math. Soc. Lecture Note Ser., vol. 440, Cambridge Univ. Press, Cambridge, 2017, pp. 312–371.
44. J. A. Noel and A. Ranganathan, *On the running time of hypergraph bootstrap percolation*, Electron. J. Combin. **30** (2023), no. 2, Paper No. 2.46, 20.
45. E. Riedl, *Largest and smallest minimal percolating sets in trees*, Electron. J. Combin. **19** (2012), no. 1, Paper 64, 18.
46. R. H. Schonmann, *On the behavior of some cellular automata related to bootstrap percolation*, Ann. Probab. **20** (1992), no. 1, 174–193.
47. A. Shapira and M. Tyomkyn, *Weakly saturated hypergraphs and a conjecture of Tuza*, Proc. Amer. Math. Soc. **151** (2023), no. 07, 2795–2805.
48. N. Terekhov, *A short proof of Tuza's conjecture for weak saturation in hypergraphs*, preprint available at arXiv:2504.03816.
49. G. L. Torrisi, M. Garetto, and E. Leonardi, *A large deviation approach to super-critical bootstrap percolation on the random graph* $G_{n,p}$, Stochastic Process. Appl. **129** (2019), no. 6, 1873–1902.
50. T. Tuza, *Asymptotic growth of sparse saturated structures is locally determined*, Discrete Math. **108** (1992), no. 1-3, 397–402.

51. S. Ulam, *Random processes and transformations*, Proceedings of the International Congress of Mathematicians, Cambridge, Mass., 1950, vol. 2, Amer. Math. Soc., Providence, RI, 1952, pp. 264–275.
52. A. C. D. van Enter, *Proof of Straley's argument for bootstrap percolation*, J. Statist. Phys. **48** (1987), no. 3-4, 943–945.
53. J. von Neumann, *Theory of self-reproducing automata*, University of Illinois Press, Champaign, IL, 1966.
54. J. T. Yu, *An extremal problem for sets: a new approach via Bezoutians*, J. Combin. Theory Ser. A **62** (1993), no. 1, 170–175.

HUN-REN ALFRÉD RÉNYI INSTITUTE OF MATHEMATICS
Email address: bartha@renyi.hu

UNIVERSITY OF WARWICK, DEPARTMENT OF STATISTICS
Email address: brett.kolesnik@warwick.ac.uk

UNIVERSITY OF OXFORD, DEPARTMENT OF MATHEMATICS
Email address: gal.kronenberg@maths.ox.ac.uk

HEBREW UNIVERSITY, EINSTEIN INSTITUTE OF MATHEMATICS
Email address: yuval.peled@mail.huji.ac.il