

DISPROOF OF THE ODD HADWIGER CONJECTURE

MARCUS KÜHN, LISA SAUERMAN, RAPHAEL STEINER, AND YUVAL WIGDERSON

ABSTRACT. We prove that there exist graphs which do not contain K_t as an odd minor and whose chromatic number is at least $(\frac{3}{2} - o(1))t$. This disproves, in a strong form, the odd Hadwiger conjecture of Gerards and Seymour from 1993.

1. INTRODUCTION

Hadwiger's conjecture [22], formulated in 1943, is one of the central open problems in graph theory. Recall that a graph F is a *minor* of a graph G if a graph isomorphic to F can be obtained from G by a sequence of vertex deletions, edge deletions, and edge contractions.

Conjecture 1.1 (Hadwiger's conjecture). *Let G be a graph. If $\chi(G) \geq t$, then G has K_t as a minor. Equivalently, if G does not have K_t as a minor, then $\chi(G) < t$.*

This is a vast generalization of the Four-Color Theorem [2, 3], which follows immediately from the $t = 5$ case of Hadwiger's conjecture. Moreover, Hadwiger's conjecture would determine the maximum chromatic number of any minor-closed family of graphs, thus unifying and generalizing a wide array of different results and conjectures bounding the chromatic number of various graph classes, see e.g. [11, 38, 43].

Despite decades of effort, Hadwiger's conjecture remains extremely far from being resolved. It is only known to hold [44] for $t \leq 6$, with the cases $t = 5$ and $t = 6$ both shown to be equivalent to the Four-Color theorem, and even the $t = 7$ case appears to be out of reach at the moment. Many researchers have attempted to prove weakenings of the conjecture, such as finding functions $f(t)$ such that every graph G not having K_t as a minor satisfies $\chi(G) < f(t)$; Hadwiger's conjecture is the assertion that this holds for $f(t) = t$. There has been a long line of works [34, 35, 40, 53, 57] steadily improving the function $f(t)$; the current record is due to Delcourt and Postle [13], who showed that this holds for $f(t) = Ct \log \log t$, where C is an absolute constant. For more on the history and progress towards Hadwiger's conjecture, we refer to the surveys [39, 45, 56].

In this paper, our focus will be on a strengthening of Hadwiger's conjecture proposed by Gerards and Seymour in 1993 (see [26, Section 6.5], [45, Section 7]), known as the *odd Hadwiger conjecture*. To explain the motivation behind this strengthening, let us go back to considering the cases of Hadwiger's conjecture when t is small. For $t = 3$, it asserts that if $\chi(G) \geq 3$, then G contains K_3 as a minor. This statement is true, since containing K_3 as a minor is equivalent to containing a cycle. But evidently, this true statement is much weaker than the "real" statement here, which is that if $\chi(G) \geq 3$, then G contains an *odd* cycle. The next case $t = 4$ of Hadwiger's conjecture is equivalent to the fact that every graph G with $\chi(G) \geq 4$ contains a K_4 -subdivision

MK: Institute for Applied Mathematics, University of Bonn, Germany. marcus.kuehn@uni-bonn.de.

LS: Institute for Applied Mathematics, University of Bonn, Germany. sauermann@iam.uni-bonn.de. Research supported by the DFG Heisenberg Program.

RS: Department of Mathematics, ETH Zürich, Switzerland. raphaelmario.steiner@math.ethz.ch. Research supported by SNSF Ambizione grant 216071.

YW: Institute for Theoretical Studies, ETH Zürich, Switzerland. yuval.wigderson@eth-its.ethz.ch. Research supported by Dr. Max Rössler, the Walter Haefner Foundation, and the ETH Zürich Foundation.

as a subgraph. And also here, a much stronger statement involving parity is known to be true: Catlin [9] proved that every such graph contains a K_4 -subdivision where every cycle formed by three subdivision paths is of odd length. Moreover, an even stronger statement is still true: Toft conjectured in 1975, and Thomassen [55] and Zang [58] independently proved, that every graph G with $\chi(G) \geq 4$ even contains a K_4 -subdivision where *all* subdivision paths have odd length. Motivated by these results and attempting to generalize Catlin's result, Gerards and Seymour proposed a natural notion, stricter than that of containing K_t as a minor, which respects the parities of cycles, and conjectured that even this stronger property is implied by large chromatic number.

More precisely, let G be a graph. We say that G has K_t as an *odd minor* if there exist vertex-disjoint trees $T_1, \dots, T_t \subseteq G$, as well as functions $\chi_i : V(T_i) \rightarrow \{\text{black}, \text{white}\}$, with the following properties:

- For every i , χ_i is a proper 2-coloring of T_i , that is, for all $\{x, y\} \in E(T_i)$, we have that $\chi_i(x) \neq \chi_i(y)$.
- For all $i \neq j$, there are vertices $x_i \in V(T_i), x_j \in V(T_j)$ such that $\{x_i, x_j\} \in E(G)$ and $\chi_i(x_i) = \chi_j(x_j)$.

With this definition in place, we can state the odd Hadwiger conjecture of Gerards and Seymour.

Conjecture 1.2 (Odd Hadwiger conjecture). *Let G be a graph. If $\chi(G) \geq t$, then G has K_t as an odd minor. Equivalently, if G does not have K_t as an odd minor, then $\chi(G) < t$.*

As discussed above, the $t = 3$ case of Conjecture 1.2 is quite straightforward, and is equivalent to the statement that every non-bipartite graph has an odd cycle. The $t = 4$ case is equivalent to the aforementioned result of Catlin. The $t = 5$ case of Conjecture 1.2 is a substantial strengthening of the Four-Color theorem, and a proof was announced by Guenin [21], but the conjecture remains wide open for $t \geq 6$.

Similarly to the case of Hadwiger's conjecture, there have been a large number of results [20, 24, 28, 29, 30, 31, 41, 42, 47, 50, 51, 54] proving weakenings of the odd Hadwiger conjecture. Interestingly, the density-based techniques used in the study of Hadwiger's conjecture do not work without serious modifications in the setting of odd minors, as they do not interact naturally with the parity requirements defining an odd minor. Despite this fact, the third author proved [49] that Hadwiger's conjecture and its odd variant are asymptotically equivalent. More precisely, he showed that, if f is any function such that any graph G not containing K_t as a minor satisfies $\chi(G) < f(t)$, then any graph G not containing K_t as an odd minor satisfies $\chi(G) < 2 \cdot f(t)$. In particular, combined with the result of Delcourt–Postle [13] mentioned above, this theorem shows that every graph G with $\chi(G) \geq Ct \log \log t$ contains K_t as an odd minor. But more importantly, this result shows that Hadwiger's conjecture and its odd variant are very closely connected; for instance, if Hadwiger's conjecture is true, then every graph G with $\chi(G) \geq 2t$ contains K_t as an odd minor.

Despite all of this evidence towards the odd Hadwiger conjecture, it turns out to be false. Our main result disproves the conjecture in the following strong form.

Theorem 1.3. *For any fixed $\delta > 0$ and every integer t which is sufficiently large with respect to δ , there is a graph G with $\chi(G) \geq (\frac{3}{2} - \delta)t$ which does not have K_t as an odd minor. In particular, for every sufficiently large t , there is a graph with $\chi(G) \geq t$ which does not have K_t as an odd minor.*

Thus, not only is the odd Hadwiger conjecture false, it is false by a factor of roughly $\frac{3}{2}$ in the sense that the chromatic number of G can exceed the order of the largest odd clique minor by an asymptotic factor of $\frac{3}{2}$. Note that, as discussed above, Hadwiger's conjecture would imply

that this ratio is at most 2. Unfortunately, as we shortly discuss further, the constant $\frac{3}{2}$ is a hard limit of our approach.

Our graph G in Theorem 1.3 forming a counterexample to the odd Hadwiger conjecture is a graph with independence number $\alpha(G) = 2$. There has been a lot of previous research on Hadwiger's conjecture and its odd variant for graphs of independence number 2. Note that if G is such a graph, then $\chi(G) \geq \lceil n/\alpha(G) \rceil = \lceil n/2 \rceil$, so Hadwiger's conjecture (resp. Conjecture 1.2) predicts that G contains a clique minor (resp. odd clique minor) of order $\lceil n/2 \rceil$. A simple argument of Duchet and Meyniel [14] shows that if G is an n -vertex graph with $\alpha(G) = 2$, then G contains a clique minor of order $\lceil n/3 \rceil$. This was extended to the odd minor setting by Kawarabayashi and Song [32], who proved that such a graph G must in fact have $K_{\lceil n/3 \rceil}$ as an odd minor. It remains a major open problem (e.g. [45, Open problem 4.8] and [10, 12]) to improve the constant factor $\frac{1}{3}$ in either of these results; see [18, Section 4] for the current best known result on Hadwiger's conjecture in graphs of independence number 2, improving the Duchet–Meyniel bound by a sublinear amount. Our counterexample to Conjecture 1.2 shows that the bound of $\lceil n/3 \rceil$ due to Kawarabayashi–Song for the odd Hadwiger conjecture in graphs of independence number 2 is asymptotically best possible. More precisely, we prove the following strengthening of Theorem 1.3.

Theorem 1.4. *For any fixed $\eta > 0$, and every integer n which is sufficiently large with respect to η , there exists an n -vertex graph G with $\alpha(G) \leq 2$ which does not have K_t as an odd minor for any $t \geq (\frac{1}{3} + \eta)n$.*

This theorem directly implies Theorem 1.3, using that any graph G with $\alpha(G) \leq 2$ satisfies $\chi(G) \geq n/\alpha(G) \geq n/2$. More details on this simple deduction can be found in Section 5.3.

The factor $\frac{1}{3}$ in Theorem 1.4 is optimal, due to the aforementioned result of Kawarabayashi and Song. We complement Theorems 1.3 and 1.4 with the following result, proved in Appendix A, showing that also the factor $\frac{3}{2}$ for the chromatic number in Theorem 1.3 is optimal for graphs G with $\alpha(G) \leq 2$.

Proposition 1.5. *For any positive integer t and any graph G with $\alpha(G) \leq 2$ which does not have K_t as an odd minor, we have $\chi(G) \leq \lceil \frac{3}{2}(t-1) \rceil$.*

2. HEURISTICS: PSEUDORANDOM TRIANGLE-FREE GRAPHS

We now turn to discuss, at a high level, the heuristics and intuition underlying our construction of the graph G in Theorem 1.4. This section can easily be skipped by readers who prefer to proceed directly to the formal construction in Section 3 and the proofs in the rest of the paper.

Since our goal is to build a graph G with $\alpha(G) \leq 2$, it is more convenient to work in the complement: we set $G = \overline{H}$, and wish to construct a triangle-free graph H whose complement does not contain K_t as an odd minor.

The simplest idea one could try is to let H be a random graph. Recall that the *binomial random graph* $G(n, p)$ is an n -vertex graph in which every unordered pair of distinct vertices forms an edge independently with probability p . As our first idea, we could sample $H \sim G(n, p)$, and aim to prove that with positive probability, we simultaneously have that H is triangle-free, and that $G = \overline{H}$ contains no large odd minor. For this discussion, we will only try to show that G does not contain $K_{n/2}$ as an odd minor—which already suffices to disprove Conjecture 1.2—as opposed to trying to obtain the optimal bound on the odd minor order given in Theorem 1.4. For simplicity, let us assume that n is even.

If G has $K_{n/2}$ as an odd minor, then there exist trees $T_1, \dots, T_{n/2}$ satisfying the definition of an odd minor. Note that since we have $n/2$ trees on n vertices, the average number of vertices of a tree T_i is at most 2. For simplicity, let us pretend for the moment that all these trees have

exactly 2 vertices (unsurprisingly, this special case is the hardest one). The structure we have in this case is called an *odd connected matching* (of size $n/2$) in the literature. Studying the presence of odd connected matchings in G , the complement of a random graph, is somewhat annoying. It will be simpler and still sufficient for us to consider an even weaker notion, which we call an *odd connected pairing*, where we do not require T_i to be a tree. More precisely, we define an *odd connected pairing of size s* in a graph G to be a set $\{(u_1, v_1), \dots, (u_s, v_s)\}$ consisting of s ordered pairs where $u_1, v_1, u_2, v_2, \dots, u_s, v_s \in V(G)$ are pairwise distinct vertices, and such that for all $i \neq j$, at least one of $\{u_i, u_j\}$ or $\{v_i, v_j\}$ is an edge of G . Note that if we further insisted that $\{u_i, v_i\}$ is an edge for all i , this would precisely correspond to having K_s as an odd minor with all the trees T_i consisting of exactly 2 vertices.

We can now begin in earnest to analyze our random construction $G = \overline{H}$, where $H \sim G(n, p)$. To get started, let us first estimate the probability that G contains an odd connected pairing of size $n/2$. For any fixed pairing $(u_1, v_1), \dots, (u_{n/2}, v_{n/2})$, and any fixed indices i, j , the probability that at least one of $\{u_i, u_j\}$ and $\{v_i, v_j\}$ is an edge of G is $1 - p^2$, since with probability p^2 both are included as edges of H . Since these events for different pairs $\{i, j\}$ of indices are independent, the probability that the fixed pairing forms an odd connected pairing in G is precisely $(1 - p^2)^{\binom{n/2}{2}}$, since we have $\binom{n/2}{2}$ pairs $\{i, j\}$ of indices. We now want to apply a union bound over all choices of the pairing; the number of such pairings equals the number of perfect matchings in K_n , which is given by $(n-1)(n-3) \cdots 1$. Asymptotically, this quantity is equal to $(n/e + o(n))^{n/2} = \exp((\frac{1}{2} + o(1))n \log n)$. Applying the union bound, we see that

$$\begin{aligned} \mathbb{P}[G \text{ contains an odd connected pairing of size } \frac{n}{2}] &\leq (n-1)(n-3) \cdots 1 \cdot (1 - p^2)^{\binom{n/2}{2}} \\ &\leq \exp\left(\left(\frac{1}{2} + o(1)\right)n \log n - p^2 \binom{n/2}{2}\right). \end{aligned}$$

In particular, if $\frac{1}{8}p^2n^2 > (\frac{1}{2} + o(1))n \log n$, then this exponent is negative for large n , and with high probability¹, G does not contain an odd connected pairing of size $n/2$. The condition $\frac{1}{8}p^2n^2 > (\frac{1}{2} + o(1))n \log n$ is equivalent to $p > (2 + o(1))\sqrt{\log n/n}$.

We also want H to be triangle-free, so that $\alpha(G) \leq 2$. Here we encounter some bad news, since it is well-known that $G(n, p)$ is with high probability triangle-free if and only if $p = o(1/n)$, much smaller than the requirement $p > (2 + o(1))\sqrt{\log n/n}$ we found above. However, we are empowered to modify H after sampling from $G(n, p)$, so long as these modifications are “insignificant” enough to not seriously affect the probability that $G = \overline{H}$ contains an odd connected pairing. It is also well-known that, so long as $p = o(1/\sqrt{n})$, a typical edge of $G(n, p)$ does not lie on any triangle; this means, intuitively, that one can delete all triangles from $G(n, p)$ without significantly affecting its structure so long as $p = o(1/\sqrt{n})$. In fact, by being very careful, we could even hope to do this for $p \leq c/\sqrt{n}$, where $c > 0$ is a small absolute constant.

Unfortunately, this is still not good enough: our requirements $p > (2 + o(1))\sqrt{\log n/n}$ and $p \leq c/\sqrt{n}$ remain incompatible. In order to explain how we overcome this, we briefly take a detour to discuss recent research on a seemingly unrelated question, which nevertheless holds the key to proving Theorem 1.4.

This seemingly unrelated question is about the asymptotics of the off-diagonal Ramsey number $r(3, k)$, which is defined as the minimum n such that every n -vertex graph contains either a triangle or an independent set of order k . Equivalently, we can ask for the minimum independence number among all n -vertex triangle-free graphs. For many decades, the best known upper bounds for this problem were obtained by a careful analysis of the random graph $G(n, c/\sqrt{n})$. Essentially, by following the intuition described above, one can carefully delete all triangles

¹Throughout the paper, we say an event occurs *with high probability* if its probability tends to 1 as $n \rightarrow \infty$.

from this random graph without significantly affecting its independence number. Different arguments along these lines were developed by Erdős [15], Spencer [48], Erdős–Suen–Winkler [16], Bollobás [6, Chapter 12], and Krivelevich [36], among others. Finally, a major breakthrough on this problem was obtained by Kim [33], who designed a different randomized procedure for generating a triangle-free graph, and proved that with high probability, this process produces an n -vertex triangle-free graph with independence number $O(\sqrt{n \log n})$, matching the classical lower bound of Ajtai–Komlós–Szemerédi [1] and Shearer [46] up to a constant factor. Subsequently, Bohman [4] analyzed a simpler random process—the so called *triangle-free process*—and proved that with high probability, it generates a pseudorandom triangle-free graph with edge density $\Omega(\sqrt{\log n/n})$. Substantially more refined results were obtained independently by Bohman and Keevash [5] and by Fiz Pontiveros, Griffiths, and Morris [17], who proved that the final edge density of the triangle-free process is $p = (\frac{1}{2} - o(1))\sqrt{\log n/n}$; moreover, they showed that in various ways, the triangle-free process is nearly indistinguishable from the binomial random graph $G(n, p)$ of the same edge density, apart from the fact that it has no triangles.

These results seem almost good enough for us: they give a pseudorandom triangle-free graph of edge density $(\frac{1}{2} - o(1))\sqrt{\log n/n}$, whereas we would like edge density $(2 + o(1))\sqrt{\log n/n}$, merely a factor of 4 away. Unfortunately, this “mere” factor of 4 is very stubborn: the results of [5, 17] also show, essentially, that it is impossible to create a denser triangle-free graph without sacrificing some of the pseudorandomness properties of the triangle-free process. This intuition led the authors of [17] to conjecture that in fact, the triangle-free process has the asymptotically smallest independence number among all n -vertex triangle-free graphs.

However, in a recent breakthrough, Campos, Jenssen, Michelen, and Sahasrabudhe [8] disproved this conjecture. One of their key insights is that for the $r(3, k)$ problem, there is no need to require the very strong pseudorandomness conditions enjoyed by the triangle-free process; one can sacrifice a small amount of pseudorandomness in order to obtain higher edge density, and this trade-off can be beneficial. This was followed by another breakthrough of Hefty, Horn, King, and Pfender [23], who found a much simpler randomized construction yielding pseudorandom triangle-free graphs of higher edge density. For their application, they select edge density $p = (1 + o(1))\sqrt{\log n/n}$ (which is expected to be optimal for the application to $r(3, k)$), but their construction naturally works at a wide range of possible edge densities. Indeed, in follow-up work on a closely related problem, Campos, Jenssen, Michelen, Pfender, and Sahasrabudhe [7] used essentially the same construction, but crucially worked with a polynomially higher edge density.

These results are the starting point of our approach, as they resolve the central incompatibility discussed above: we now have access to a construction of pseudorandom triangle-free graphs whose edge density can be larger than the threshold $(2 + o(1))\sqrt{\log n/n}$ we encountered above. In fact, in order to simplify the proofs of our main technical results, it will be convenient to work with a logarithmically higher edge density, namely $p = \Theta((\log n)^2/\sqrt{n})$.

The basic idea in the works [7, 8, 23] is the following simple approach. Let H_{start} be a triangle-free graph on m vertices, for some m that is smaller but not much smaller than n , and let H_{blowup} be obtained from H_{start} by blowing up each vertex by a factor of n/m . That is, we replace each vertex of H_{start} by an independent set of order n/m , and each edge of H_{start} by a complete bipartite graph between the independent sets corresponding to its endpoints. Then on the one hand, if H_{start} is pseudorandom in an appropriate sense, then H_{blowup} retains many of the pseudorandomness properties; for example, the edge distribution among all large sets in H_{blowup} is roughly the same as it was in H_{start} . Moreover, the two graphs $H_{\text{blowup}}, H_{\text{start}}$ have essentially the same edge density. However, when viewed as a function of n , the edge density of H_{blowup} is higher than that of H_{start} . For example, if H_{start} had $m^{3/2}$ edges, then H_{blowup}

has $m^{3/2} \cdot (n/m)^2 = n^2/\sqrt{m}$ edges, and, if $m = o(n)$, then $n^2/\sqrt{m} = \omega(n^{3/2})$. In particular, if we take $H_{\text{start}} \sim G(m, m^{-1/2})$, then H_{blowup} will be an n -vertex graph with edge density higher than $n^{-1/2}$, which still maintains some of the pseudorandomness properties of $G(m, m^{-1/2})$.

Of course, for many applications, including the application to $r(3, k)$ and to the odd Hadwiger conjecture, we have done essentially nothing: although we have artificially increased the relative edge density, the large independent sets in H_{blowup} completely counteract whatever we gained from this. The other new idea in [7, 8, 23] is to “conceal” the blowup structure of H_{start} by overlaying it with another, independent source of (pseudo)randomness. In [8] this is accomplished by running the triangle-free process on the blown-up graph H_{blowup} , but the much simpler approach taken in [7, 23] is simply to randomly overlay two blowups. Because the overlaying is done in a random way, the blowup structures of the two graphs are essentially “orthogonal”, and the edges of one blowup destroy the structure of the other blowup. This is also the approach we take.

More specifically, we start with two binomial random graphs with distribution $G(m, p)$, where $p = o(1/\sqrt{m})$, but $p = \omega(\sqrt{\log n/n})$. To make them triangle-free, in each of these graphs we first delete all edges occurring in triangles. Since $p = o(1/\sqrt{m})$, we expect a negligible proportion of these edges to be deleted this way. Thus, we expect each of the resulting random graphs to have roughly $pm^2/2$ edges. We now blow each of these graphs up by a factor of n/m and randomly overlay these blowups (for more details on this step, see the next section). This way, we obtain an n -vertex graph with roughly $2 \cdot (pm^2/2) \cdot (n/m)^2 = pn^2$ edges. Again, we expect few edges to participate in triangles consisting of edges from both of the blow-up graphs. We then delete an edge from every such triangle in a certain way, so that at the end of this process we expect to end up with a triangle-free graph H of edge density roughly p ; crucially, we have that $p = \omega(\sqrt{\log n/n})$, and thus our heuristic from the binomial random graph suggests that \overline{H} is unlikely to contain an odd connected pairing of size $n/2$.

3. THE CONSTRUCTION

Let n be a large integer. We fix² parameters

n, m, p

$$m = \left\lceil \frac{n}{(\log n)^8} \right\rceil \quad \text{and} \quad p = \frac{1}{m^{1/2} \cdot (\log m)^2}.$$

Our construction is based on the following independent random choices:

- a binomial random graph $H_R \sim G(m, p)$ with vertex set $V_R = [m]$, H_R
- a binomial random graph $H_B \sim G(m, p)$ with vertex set $V_B = [m]$, and H_B
- two independent, uniformly random functions $\pi_R : [n] \rightarrow [m]$ and $\pi_B : [n] \rightarrow [m]$. Here, π_R, π_B
we think of π_R as a function to V_R and π_B as a function to V_B .

We now let H'_R be obtained from H_R by deleting every edge that lies in a triangle, and similarly for H'_B . We then define a graph H_0 with vertex set $[n]$ by taking the union of the pull-backs of H'_R and H'_B along the maps π_R and π_B , respectively. Formally, we join distinct $u, v \in [n]$ by an edge of H_0 if

H'_R, H'_B, H_0

$$\{\pi_R(u), \pi_R(v)\} \in E(H'_R) \quad \text{or} \quad \{\pi_B(u), \pi_B(v)\} \in E(H'_B).$$

For an edge $\{u, v\} \in E(H_0)$, we call the edge $\{u, v\}$ *red* if $\{\pi_R(u), \pi_R(v)\} \in E(H'_R)$, and we call the edge $\{u, v\}$ *blue* if $\{\pi_B(u), \pi_B(v)\} \in E(H'_B)$. Note that the graph of red edges is precisely a blowup of H'_R , where vertex $x \in V_R$ is replaced by an independent set of size $|\pi_R^{-1}(x)|$, and similarly for the blue edges. Also note that some edges may receive both colors. On the other

²We remark that there is a lot of flexibility in these choices; we really only need that n/m is at least some small fixed power of $\log n$ and at most some small fixed power of n , and that the same holds for $p^2 n$ and $(p^2 m)^{-1}$.

hand, since H'_R, H'_B are triangle-free graphs, there is no triangle in H_0 all of whose edges receive the same color.

Thus, every triangle in H_0 must comprise two red edges and one blue edge, or two blue edges and one red edge. We then define H to be the subgraph of H_0 obtained by deleting the edge of the minority color from each triangle. That is, for every red edge $\{u, v\} \in E(H_0)$ for which there exists w such that $\{u, w\}, \{v, w\} \in E(H_0)$ are blue edges, we delete $\{u, v\}$ from H_0 in constructing H ; and similarly for every blue edge $\{u, v\} \in E(H_0)$ for which there exists w such that $\{u, w\}, \{v, w\} \in E(H_0)$ are red edges. It is not relevant for what follows, but it may be helpful to note that every deleted edge has only one color, since there are no monochromatic triangles in H_0 .

At the end of this process, we have constructed an n -vertex triangle-free graph H . We conclude the construction by taking $G = \overline{H}$ to be the complement of H , so that $\alpha(G) \leq 2$. As mentioned in the previous section, an *odd connected pairing of size s* in a graph G is defined to be a set $\{(u_1, v_1), \dots, (u_s, v_s)\}$ consisting of s ordered pairs where $u_1, v_1, u_2, v_2, \dots, u_s, v_s \in V(G)$ are pairwise distinct vertices, and such that for all $i \neq j$, at least one of $\{u_i, u_j\}$ or $\{v_i, v_j\}$ is an edge of G . We will prove that with high probability, G contains no odd connected pairing of size at least εn , which then implies the following theorem.

Theorem 3.1. *For every fixed $\varepsilon > 0$ and all sufficiently large n , there exists an n -vertex graph G with $\alpha(G) \leq 2$ which does not contain an odd connected pairing of size at least εn .*

It is not hard to show, as we do in Section 5.3, that Theorem 3.1 implies Theorem 1.4, and hence Theorem 1.3. Thus, for the majority of the rest of the paper, we focus on proving Theorem 3.1.

Organization and notation. In Section 4, we state the main probabilistic lemmas we will need to analyze our construction. We then prove Theorem 3.1 using these lemmas in Section 5 and afterwards deduce Theorems 1.3 and 1.4. We prove the probabilistic lemmas in Sections 6 and 7, and end in Section 8 with some concluding remarks. Finally, we prove Proposition 1.5 in Appendix A.

For a non-negative integer n , we write $[n] := \{1, \dots, n\}$, as usual. All logarithms are to base e . To save parentheses, we write $\log^k n$ instead of $(\log n)^k$. Given a set S and an integer k , we use $\binom{S}{k}$ to denote the collection of all k -element subsets of S . For an event \mathcal{E} , we write \mathcal{E}^c to denote the complementary event. We furthermore use the standard asymptotic notation O, Ω, o , and ω . For a graph G , we denote by Δ_G its maximum degree, by $e(G)$ its number of edges, by $\chi(G)$ its chromatic number, by $\alpha(G)$ its independence number, and by $\nu(G)$ its matching number.

In order to help the reader track the main definitions, we include marginal notes to indicate the first time a key notion is introduced.

4. PROBABILISTIC LEMMAS

In order to prove Theorem 3.1, we will analyze a number of properties depending on the random choices of H_R, H_B, π_R, π_B . In this section, we introduce various events defined by these random choices, and state a number of lemmas which say, roughly, that various “good” outcomes happen with high probability. In the next section, we use these probabilistic lemmas to prove Theorem 3.1. We now introduce each of these events in turn.

The first four are quite straightforward: they say that the random graphs H_R, H_B and the random functions π_R, π_B are “balanced”, in the sense that the graphs do not have abnormally high maximum degree, and the functions do not have abnormally large fibers. Formally, we define \mathcal{D}_R to be the event that the maximum degree of H_R is at most $2pm$, and similarly \mathcal{D}_B to be the event that the maximum degree of H_B is at most $2pm$. We also define \mathcal{F}_R to be the event

 H

odd connected pairing

 $\mathcal{D}_R, \mathcal{D}_B$

that $|\pi_R^{-1}(x)| \leq 2n/m$ for all $x \in V_R$, and similarly \mathcal{F}_B to be the event that $|\pi_B^{-1}(x)| \leq 2n/m$ for all $x \in V_B$. Note that the events $\mathcal{D}_R, \mathcal{D}_B, \mathcal{F}_R$ and \mathcal{F}_B only depend on the outcome of the random object H_R, H_B, π_R , and π_B , respectively. The following simple lemma, an elementary consequence of the Chernoff bound, shows that these four events hold with high probability (the proof can be found in Section 7.3). $\mathcal{F}_R, \mathcal{F}_B$

Lemma 4.1. *The events $\mathcal{D}_R, \mathcal{D}_B, \mathcal{F}_R, \mathcal{F}_B$ hold with high probability as $n \rightarrow \infty$.*

In our next condition, we aim to control the interactions of the two functions π_R and π_B . The point here is that these maps should be “orthogonal”, so that the randomness from H_R and from H_B interact a great deal in defining H . However, the formal condition requires some setup. For $s \in \mathbb{N}$, an *s-pairing* in $[n]$ is a collection of pairs $\{(u_1, v_1), \dots, (u_s, v_s)\}$, where $u_1, \dots, u_s, v_1, \dots, v_s$ are distinct elements of $[n]$. For an *s-pairing* $P = \{(u_1, v_1), \dots, (u_s, v_s)\}$ in $[n]$, and $\varepsilon > 0$, we say that a map $\pi : [n] \rightarrow [m]$ is ε -*respectful* of P if there is a subset $I \subseteq [s]$ with $|I| \geq \varepsilon s$ such that $\pi(u_i) \neq \pi(v_i)$ for all $i \in I$, and furthermore the pairs $\{\pi(u_i), \pi(v_i)\}$ are distinct for all $i \in I$. That is, π is ε -respectful of the pairing P if at least εs of the pairs in P get mapped to distinct pairs in $\binom{[m]}{2}$. Finally, for $\varepsilon > 0$, let $\mathcal{E}(\varepsilon)$ be the event that for each *s-pairing* P in $[n]$ with $s = \lceil \varepsilon n \rceil$, at least one of the random functions π_R and π_B is ε -respectful of P . Note that this event only depends on the random functions π_R and π_B , but not on the random graphs H_R and H_B . The following lemma, proved in Section 6, shows that the event $\mathcal{E}(\varepsilon)$ happens with high probability. ε -respectful
 $\mathcal{E}(\varepsilon)$

Lemma 4.2. *Fix $0 < \varepsilon < \frac{1}{10}$. Then the event $\mathcal{E}(\varepsilon)$ holds with high probability as $n \rightarrow \infty$.*

Recall that the graph H'_R is obtained from H_R by deleting all edges in triangles, and H'_B is obtained analogously from H_B . Note that in H_R , the expected number of triangles is much less than the expected number of edges, hence a typical edge is unlikely to be on a triangle. As a consequence, we expect few edges to be deleted when passing from H_R to H'_R . However, for different edges, there are also dependencies in the probabilities of these edges being on triangles and hence being deleted when passing from H_R to H'_R . In our probabilistic analysis, we therefore first consider a somewhat denser random graph H_R^* , and then sample H_R by randomly sparsifying H_R^* , selecting edges of H_R^* with an appropriate probability. Crucially, in this intermediate random graph H_R^* , with high probability every edge is contained in only very few triangles. This is helpful, since we then have much fewer dependencies for different edges to lie on triangles, when further sparsifying H_R^* to obtain H_R .

Formally, we consider independent binomial random graphs $H_R^* \sim G(m, m^{-1/2})$ with vertex set $V_R = [m]$ and $H_B^* \sim G(m, m^{-1/2})$ with vertex set $V_B = [m]$ (also independent from π_R and π_B), where H_R^* is coupled with H_R such that $H_R \subseteq H_R^*$ always holds, and H_B^* is coupled with H_B such that $H_B \subseteq H_B^*$ always holds. We furthermore choose these couplings in such a way that, conditioning on an outcome \widehat{H}_R^* of H_R^* , the distribution of the random graph H_R is given by taking every edge of \widehat{H}_R^* into H_R independently at random with probability $p/m^{-1/2} = 1/\log^2 m$ (also independent from all other random choices related to π_R, π_B, H_B^* , and H_B); and analogously for H_B when conditioning on an outcome \widehat{H}_B^* of H_B^* . H_R^*, H_B^*

Now, let \mathcal{T}_R be the event that every edge $e \in E(H_R^*)$ is contained in at most $20 \log m$ triangles in H_R^* . Analogously, let \mathcal{T}_B be the event that every edge $e \in E(H_B^*)$ is contained in at most $20 \log m$ triangles in H_B^* . It is not hard to show that these events \mathcal{T}_R and \mathcal{T}_B hold with high probability; this again follows from an easy Chernoff bound argument (see Section 7.3 for the details). $\mathcal{T}_R, \mathcal{T}_B$

Lemma 4.3. *The events \mathcal{T}_R and \mathcal{T}_B hold with high probability as $n \rightarrow \infty$.*

As long as the event \mathcal{T}_R holds, we have sufficiently good control on the dependencies between different edges to end up in triangles in H_R and thus getting deleted when passing to H'_R . Specifically, for a list of pairs of edges of K_m (corresponding to edge pairs of the form $(\{\pi_R(u_i), \pi_R(u_j)\}, \{\pi_R(v_i), \pi_R(v_j)\})$ for $i \neq j$ in some potential odd connected pairing), we are interested in the probability that for at least one edge pair in the list both members are present as edges in H_R and are both not on triangles in H_R . If this happens, both members will be present as edges in H'_R and therefore lead to edges in H_0 . Our next lemma shows that on the event \mathcal{T}_R , we are indeed quite likely to find one such edge pair in the list, with a probability estimate that is strong enough to union-bound over all potential odd connected pairings.

Lemma 4.4. *Fix $X \in \{R, B\}$ and $\gamma > 0$, and assume that n is sufficiently large with respect to γ . Consider a graph Γ with vertex set $\binom{[n]}{2}$ such that $e(\Gamma) \geq \gamma n^2$ and $\Delta_\Gamma \leq 4n^2/m^2$. Then we have*

$$\mathbb{P}[\mathcal{T}_X \text{ holds and there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H'_X)] \leq n^{-2n}.$$

Note that for $X = R$, the event whose probability is considered in this lemma only depends on the outcomes of H_R^* and H_R (and not on the outcomes of π_R, π_B, H_B^* and H_B), and similarly for $X = B$ the event only depends on the outcomes of H_B^* and H_B .

5. PROOF OF THE MAIN THEOREMS

In this section, we prove Theorems 1.3 and 1.4 assuming the probabilistic lemmas stated in the previous section. We will need two more combinatorial lemmas, which we prove in the first subsection. The second subsection contains the proof of Theorem 3.1 (assuming the probabilistic lemmas from Section 4), from which we then finally deduce Theorems 1.3 and 1.4 in the last subsection.

5.1. Combinatorial lemmas. Let us say that a pair $\{u, v\}$ of vertices $u, v \in [n]$ is *closed by a red cherry* if there is a vertex $w \in [n]$ so that $\{u, w\}, \{v, w\} \in E(H_0)$ and both these edges are red. We define what it means to be *closed by a blue cherry* analogously. The pairs closed by red or blue cherries are precisely those pairs that we would be forced to delete in the creation of H from H_0 : if a pair $\{u, v\}$ closed by a red or blue cherry is included as an edge of H_0 , it will not survive into H .

closed by red or
blue cherry

Lemma 5.1. *Assume that n is sufficiently large. Then, whenever the events \mathcal{D}_R and \mathcal{F}_R hold, the number of pairs $\{u, v\} \in \binom{[n]}{2}$ closed by a red cherry in H_0 is at most $n^2/\log n$. Similarly, whenever the events \mathcal{D}_B and \mathcal{F}_B hold, the number of pairs $\{u, v\} \in \binom{[n]}{2}$ closed by a blue cherry in H_0 is at most $n^2/\log n$.*

Proof. By symmetry it suffices to prove the first statement, regarding red cherries. If $\{u, v\} \in \binom{[n]}{2}$ is closed by a red cherry, this means that there exists w such that $\{u, w\}, \{v, w\}$ are red edges of H_0 . This, in turn, means that $\{\pi_R(u), \pi_R(w)\}, \{\pi_R(v), \pi_R(w)\}$ are edges of H'_R and hence of H_R , so $\pi_R(u)$ and $\pi_R(v)$ must have a common neighbor in the graph H_R . Since \mathcal{D}_R holds, all vertices of H_R have degree at most $2pm$, and so the number of pairs of vertices of H_R with a common neighbor is at most $m \cdot (2pm)^2 = 4p^2m^3$. Furthermore, for every pair of vertices $x, y \in V_R$ with a common neighbor in H_R , there are only $|\pi_R^{-1}(x)| \leq 2n/m$ options for $u \in [n]$ with $\pi_R(u) = x$ and only $|\pi_R^{-1}(y)| \leq 2n/m$ options for $v \in [n]$ with $\pi_R(v) = y$, since the event \mathcal{F}_R holds. This means that the total number of pairs $\{u, v\} \in \binom{[n]}{2}$ is closed by a red cherry is at most $4p^2m^3 \cdot (2n/m)^2 = 16p^2mn^2 = 16n^2/\log^4 m \leq n^2/\log n$ (recalling that $m = \lceil n/\log^8 n \rceil$, $p = m^{-1/2} \log^{-2} m$, and n is sufficiently large). \square

Recall that our goal is to prove Theorem 3.1. So we want to show that the graph G we constructed is unlikely to contain an odd connected pairing of size at least εn . Recall that any given s -pairing $P = \{(u_1, v_1), \dots, (u_s, v_s)\}$ in $[n]$ with $s \geq \varepsilon n$ forms an odd connected pairing in G if and only if for all distinct $i, j \in [s]$, at least one of $\{u_i, u_j\}$ and $\{v_i, v_j\}$ is an edge of G . To understand whether the pairs $\{u_i, u_j\}$ and $\{v_i, v_j\}$ form edges in G , we need to consider their projections $\{\pi_R(u_i), \pi_R(u_j)\}$ and $\{\pi_R(v_i), \pi_R(v_j)\}$ to the vertex set of H_R (and similarly for H_B). Our goal is to analyze the probability that for some $i, j \in [s]$ both $\{\pi_R(u_i), \pi_R(u_j)\}$ and $\{\pi_R(v_i), \pi_R(v_j)\}$ form edges of H_R , or more precisely, of H'_R (because then $\{u_i, u_j\}$ and $\{v_i, v_j\}$ will be edges of H_0). For this analysis, the first important step is to show that a lot of pairs of vertex pairs of H_R appear as $\{\{\pi_R(u_i), \pi_R(u_j)\}, \{\pi_R(v_i), \pi_R(v_j)\}\}$ for some $i, j \in [s]$. The following lemma shows this, assuming that π_R is ε -respectful of P and satisfies \mathcal{F}_R .

Lemma 5.2. *Let $X \in \{R, B\}$, fix $0 < \varepsilon < 1$, and let n be sufficiently large with respect to ε . Let $P = \{(u_1, v_1), \dots, (u_s, v_s)\}$ be an s -pairing in $[n]$ such that $s \geq \varepsilon n$. Consider an outcome $\widehat{\pi}_X$ of the random map $\pi_X : [n] \rightarrow [m]$ which is ε -respectful of P and such that the event \mathcal{F}_X holds for $\pi_X = \widehat{\pi}_X$.*

Let Π be the set of all pairs $\{\{\widehat{\pi}_X(u_i), \widehat{\pi}_X(u_j)\}, \{\widehat{\pi}_X(v_i), \widehat{\pi}_X(v_j)\}\}$ for all $i, j \in [s]$ with $\widehat{\pi}_X(u_i) \neq \widehat{\pi}_X(u_j)$ and $\widehat{\pi}_X(v_i) \neq \widehat{\pi}_X(v_j)$ and also $\{\widehat{\pi}_X(u_i), \widehat{\pi}_X(u_j)\} \neq \{\widehat{\pi}_X(v_i), \widehat{\pi}_X(v_j)\}$. Then we have $|\Pi| \geq \varepsilon^4 n^2 / 6$.

We can imagine Π in Lemma 5.2 as a graph on the vertex set $\binom{[m]}{2}$. In this graph, for any $i, j \in [s]$, we draw an edge between $\{\widehat{\pi}_X(u_i), \widehat{\pi}_X(u_j)\}$ and $\{\widehat{\pi}_X(v_i), \widehat{\pi}_X(v_j)\}$, under the conditions that $\widehat{\pi}_X(u_i) \neq \widehat{\pi}_X(u_j)$ (as otherwise $\{\widehat{\pi}_X(u_i), \widehat{\pi}_X(u_j)\}$ does not correspond to an element of $\binom{[m]}{2}$), and similarly $\widehat{\pi}_X(v_i) \neq \widehat{\pi}_X(v_j)$ as well as $\{\widehat{\pi}_X(u_i), \widehat{\pi}_X(u_j)\} \neq \{\widehat{\pi}_X(v_i), \widehat{\pi}_X(v_j)\}$ (since we do not want to create loops in our graph on $\binom{[m]}{2}$). We do not draw parallel edges, so if some pair of vertices in $\binom{[m]}{2}$ appears as $\{\{\widehat{\pi}_X(u_i), \widehat{\pi}_X(u_j)\}, \{\widehat{\pi}_X(v_i), \widehat{\pi}_X(v_j)\}\}$ for multiple different choices of $i, j \in [s]$, we still only draw a single edge. The conclusion of the lemma is that this graph has at least $\varepsilon^4 n^2 / 6$ edges.

Proof of Lemma 5.2. Without loss of generality let $X = R$ (the case $X = B$ is analogous). Exclusively for this proof, we adopt the notation \bar{v} to denote $\widehat{\pi}_R(v)$. By the assumption that $\widehat{\pi}_R : [n] \rightarrow [m]$ is ε -respectful of P , there exists a subset $I \subseteq [s]$ of size $|I| \geq \varepsilon s \geq \varepsilon^2 n$ such that $\bar{u}_i \neq \bar{v}_i$ for all $i \in I$, and furthermore the pairs $\{\bar{u}_i, \bar{v}_i\}$ for $i \in I$ are pairwise distinct.

Note that for any $i, j \in I$ with $\bar{u}_i \neq \bar{u}_j$ and $\bar{v}_i \neq \bar{v}_j$, we also have $\{\bar{u}_i, \bar{u}_j\} \neq \{\bar{v}_i, \bar{v}_j\}$. Indeed, if $\{\bar{u}_i, \bar{u}_j\} = \{\bar{v}_i, \bar{v}_j\}$, then due to $\bar{u}_i \neq \bar{v}_i$ and $\bar{u}_j \neq \bar{v}_j$, we would have $\bar{u}_i = \bar{v}_j$ and $\bar{u}_j = \bar{v}_i$ and hence $\{\bar{u}_i, \bar{v}_i\} = \{\bar{u}_j, \bar{v}_j\}$ for the distinct elements $i, j \in I$ (we must have $i \neq j$ because of $\bar{u}_i \neq \bar{u}_j$). This would be a contradiction.

Recall that Π is the set of all pairs $\{\{\bar{u}_i, \bar{u}_j\}, \{\bar{v}_i, \bar{v}_j\}\}$ for all $i, j \in [s]$ with $\bar{u}_i \neq \bar{u}_j$ and $\bar{v}_i \neq \bar{v}_j$ and also $\{\bar{u}_i, \bar{u}_j\} \neq \{\bar{v}_i, \bar{v}_j\}$. Now, let us just consider the pairs $\{\{\bar{u}_i, \bar{u}_j\}, \{\bar{v}_i, \bar{v}_j\}\}$ for all $i, j \in I$ with $\bar{u}_i \neq \bar{u}_j$ and $\bar{v}_i \neq \bar{v}_j$ (recall that we then also have $\{\bar{u}_i, \bar{u}_j\} \neq \{\bar{v}_i, \bar{v}_j\}$). Together, these pairs form some subset $\Pi_I \subseteq \Pi$, and it suffices to show $|\Pi_I| \geq \varepsilon^4 n^2 / 6$.

The key claim is that for any given $x, y, z, w \in [m]$, the pair $\{\{x, y\}, \{z, w\}\}$ can appear as $\{\{\bar{u}_i, \bar{u}_j\}, \{\bar{v}_i, \bar{v}_j\}\}$ for at most two different choices of $\{i, j\} \subseteq I$. Indeed, if

$$\{\{x, y\}, \{z, w\}\} = \{\{\bar{u}_i, \bar{u}_j\}, \{\bar{v}_i, \bar{v}_j\}\},$$

then the two pairs $\{\bar{u}_i, \bar{v}_i\}$ and $\{\bar{u}_j, \bar{v}_j\}$ must be either the two pairs $\{x, z\}$ and $\{y, w\}$ or the two pairs $\{x, w\}$ and $\{y, z\}$. More precisely, $\{\{\bar{u}_i, \bar{v}_i\}, \{\bar{u}_j, \bar{v}_j\}\}$ must be either $\{\{x, z\}, \{y, w\}\}$ or $\{\{x, w\}, \{y, z\}\}$. Since the pairs $\{\bar{u}_i, \bar{v}_i\}$ are distinct for all $i \in I$, in either of these two cases there is at most one possibility for $\{i, j\} \subseteq I$. Thus, in total, there are indeed at most two choices for $\{i, j\} \subseteq I$.

Consequently, for any $\{i, j\} \subseteq I$ with $\overline{u_i} \neq \overline{u_j}$ and $\overline{v_i} \neq \overline{v_j}$, the pair $\{\{\overline{u_i}, \overline{u_j}\}, \{\overline{v_i}, \overline{v_j}\}\}$ is an element of Π_I , and every element of Π_I is obtained in this way for at most two choices of $\{i, j\} \subseteq I$. Therefore we have

$$|\Pi_I| \geq \frac{|\{\{i, j\} \subseteq I : \overline{u_i} \neq \overline{u_j}, \overline{v_i} \neq \overline{v_j}\}|}{2}.$$

Among the total number $\binom{|I|}{2}$ of pairs $\{i, j\} \subseteq I$, the number of pairs with $\overline{u_i} = \overline{u_j}$ is at most $|I| \cdot (2n/m)$. Indeed, choosing any $j \in I$, in order to have $\overline{u_i} = \overline{u_j}$, we need to choose $i \in I$ such that u_i is a vertex in the set $(\widehat{\pi_R})^{-1}(x)$ for $x = \overline{u_j}$. Due to the event \mathcal{F}_R we have $|\widehat{\pi_R}^{-1}(x)| \leq 2n/m$, so there are at most $2n/m$ choices for u_i and hence for $i \in I$ (recall that the vertices $u_1, \dots, u_s, v_1, \dots, v_s$ are all distinct). Similarly, the number of pairs $\{i, j\} \subseteq I$ with $\overline{v_i} = \overline{v_j}$ is at most $|I| \cdot (2n/m)$. Thus,

$$\begin{aligned} |\{\{i, j\} \subseteq I : \overline{u_i} \neq \overline{u_j}, \overline{v_i} \neq \overline{v_j}\}| &\geq \binom{|I|}{2} - |I| \cdot \frac{2n}{m} - |I| \cdot \frac{2n}{m} \\ &= |I| \left(\frac{|I| - 1}{2} - \frac{4n}{m} \right) \\ &\geq \varepsilon^2 n \left(\frac{\varepsilon^2 n - 1}{2} - \frac{4n}{m} \right) \geq \frac{\varepsilon^4 n^2}{3}, \end{aligned}$$

using that n and hence m is sufficiently large with respect to ε . This implies $|\Pi| \geq |\Pi_I| \geq \varepsilon^4 n^2 / 6$, as desired. \square

5.2. Proof of Theorem 3.1. Given the results of the previous subsection, we are now ready for the proof of Theorem 3.1.

Proof of Theorem 3.1, assuming the lemmas in Section 4. Note that we may assume without loss of generality that the given fixed value of ε satisfies $0 < \varepsilon < 1/10$ (since the statement in Theorem 3.1 for some value of ε automatically implies the same statement for all larger values of ε). Let us define $\gamma := \varepsilon^4/10 > 0$. We are assuming that n is sufficiently large with respect to ε , so n is also large with respect to γ (in particular, large enough for the statement in Lemma 4.4 to hold).

Recall that the random graph G constructed in Section 3 has n vertices and satisfies $\alpha(G) \leq 2$. Our goal is to show that, as $n \rightarrow \infty$, with high probability G does not have an odd connected pairing of size at least εn . This in particular implies that for sufficiently large n (sufficiently large with respect to ε), there exists an n -vertex graph G with $\alpha(G) \leq 2$ which does not contain an odd connected pairing of size at least εn .

Letting $s = \lceil \varepsilon n \rceil$, observe that G has an odd connected pairing of size at least εn if and only if it has an odd connected pairing of size s , meaning that some s -pairing in $[n]$ forms an odd connected pairing in G . So we can observe that

$$\begin{aligned} &\mathbb{P}[G \text{ has an odd connected pairing of size at least } \varepsilon n] \\ &\leq \mathbb{P}[\mathcal{F}_R^c] + \mathbb{P}[\mathcal{F}_B^c] + \mathbb{P}[\mathcal{D}_R^c] + \mathbb{P}[\mathcal{D}_B^c] + \mathbb{P}[\mathcal{T}_R^c] + \mathbb{P}[\mathcal{T}_B^c] + \mathbb{P}[\mathcal{E}(\varepsilon)^c] \\ &\quad + \mathbb{P}[\mathcal{F}_R, \mathcal{F}_B, \mathcal{D}_R, \mathcal{D}_B, \mathcal{T}_R, \mathcal{T}_B, \mathcal{E}(\varepsilon) \text{ hold and } G \text{ has an odd connected pairing of size } s] \\ &\leq \mathbb{P}[\mathcal{F}_R^c] + \mathbb{P}[\mathcal{F}_B^c] + \mathbb{P}[\mathcal{D}_R^c] + \mathbb{P}[\mathcal{D}_B^c] + \mathbb{P}[\mathcal{T}_R^c] + \mathbb{P}[\mathcal{T}_B^c] + \mathbb{P}[\mathcal{E}(\varepsilon)^c] \\ &\quad + \sum_P \mathbb{P}[\mathcal{F}_R, \mathcal{F}_B, \mathcal{D}_R, \mathcal{D}_B, \mathcal{T}_R, \mathcal{T}_B, \mathcal{E}(\varepsilon) \text{ hold and } P \text{ is an odd connected pairing in } G], \end{aligned}$$

where the sum is taken over all s -pairings P in $[n]$. Note that the number of such s -pairings is at most n^n . Thus, it suffices to prove

$$\mathbb{P}[\mathcal{F}_R, \mathcal{F}_B, \mathcal{D}_R, \mathcal{D}_B, \mathcal{T}_R, \mathcal{T}_B, \mathcal{E}(\varepsilon) \text{ hold and } P \text{ is an odd connected pairing in } G] \leq n^{-2n} \quad (1)$$

for every s -pairing P in $[n]$. Indeed, then using Lemmas 4.1–4.3, we can conclude that

$$\begin{aligned} & \mathbb{P}[G \text{ has an odd connected pairing of size at least } \varepsilon n] \\ & \leq o(1) + n^n \cdot n^{-2n} \\ & = o(1) + n^{-n} = o(1), \end{aligned}$$

as desired.

To prove (1) consider an s -pairing $P = \{(u_1, v_1), \dots, (u_s, v_s)\}$ in $[n]$. It suffices to prove a probability bound as in (1) conditional on any outcomes $\widehat{\pi}_R$ and $\widehat{\pi}_B$ of the random maps π_R and π_B . If one of the events \mathcal{F}_R , \mathcal{F}_B and $\mathcal{E}(\varepsilon)$ fails for these outcomes $\widehat{\pi}_R$ and $\widehat{\pi}_B$, the relevant conditional probability is zero. So we may assume that \mathcal{F}_R , \mathcal{F}_B and $\mathcal{E}(\varepsilon)$ hold when $\pi_R = \widehat{\pi}_R$ and $\pi_B = \widehat{\pi}_B$. Then, by definition of the event $\mathcal{E}(\varepsilon)$, at least one of the functions $\widehat{\pi}_R$ and $\widehat{\pi}_B$ must be ε -respectful of P . Let us assume without loss of generality that $\widehat{\pi}_R$ is ε -respectful of P (the other case is analogous).

We can now apply Lemma 5.2, since $\widehat{\pi}_R : [n] \rightarrow [m]$ is ε -respectful of the s -pairing $P = \{(u_1, v_1), \dots, (u_s, v_s)\}$ (with $s \geq \varepsilon n$), and \mathcal{F}_R holds for $\pi_R = \widehat{\pi}_R$. As in the lemma statement, let Π be the set of all pairs $\{\{\widehat{\pi}_R(u_i), \widehat{\pi}_R(u_j)\}, \{\widehat{\pi}_R(v_i), \widehat{\pi}_R(v_j)\}\}$ for all $i, j \in [s]$ with $\widehat{\pi}_R(u_i) \neq \widehat{\pi}_R(u_j)$ and $\widehat{\pi}_R(v_i) \neq \widehat{\pi}_R(v_j)$ and also $\{\widehat{\pi}_R(u_i), \widehat{\pi}_R(u_j)\} \neq \{\widehat{\pi}_R(v_i), \widehat{\pi}_R(v_j)\}$. Then by Lemma 5.2 we have $|\Pi| \geq \varepsilon^4 n^2 / 6$.

Let us now also condition on any fixed outcome \widehat{H}_B of the random graph H_B . If the event \mathcal{D}_B fails for $H_B = \widehat{H}_B$, the relevant probability we want to bound is again zero. So we may assume that \mathcal{D}_B holds for $H_B = \widehat{H}_B$. Thus, we can apply Lemma 5.1, and conclude that at most $n^2 / \log n$ pairs $\{u, v\} \in \binom{[n]}{2}$ are closed by a blue cherry in H_0 whenever $H_B = \widehat{H}_B$ and $\pi_B = \widehat{\pi}_B$. Note that the outcomes of π_B and H_B completely determine the blue edges in H_0 , and they therefore determine which pairs of vertices are closed by a blue cherry in H_0 .

Let $\Pi_\Gamma \subseteq \Pi$ be the set of all pairs $\{\{\widehat{\pi}_R(u_i), \widehat{\pi}_R(u_j)\}, \{\widehat{\pi}_R(v_i), \widehat{\pi}_R(v_j)\}\}$ for all $i, j \in [s]$ with $\widehat{\pi}_R(u_i) \neq \widehat{\pi}_R(u_j)$ and $\widehat{\pi}_R(v_i) \neq \widehat{\pi}_R(v_j)$ and $\{\widehat{\pi}_R(u_i), \widehat{\pi}_R(u_j)\} \neq \{\widehat{\pi}_R(v_i), \widehat{\pi}_R(v_j)\}$, such that neither $\{u_i, u_j\}$ nor $\{v_i, v_j\}$ is closed by a blue cherry in H_0 when $H_B = \widehat{H}_B$ and $\pi_B = \widehat{\pi}_B$. Since the number of $\{i, j\} \subseteq [s]$ such that $\{u_i, u_j\}$ or $\{v_i, v_j\}$ is closed by such a blue cherry is at most $n^2 / \log n$, we find that

$$|\Pi_\Gamma| \geq |\Pi| - \frac{n^2}{\log n} \geq \frac{\varepsilon^4 n^2}{6} - \frac{n^2}{\log n} \geq \gamma n^2,$$

recalling that $\gamma = \varepsilon^4 / 10$ and using the assumption that n is sufficiently large with respect to ε . Furthermore, the set Π_Γ is determined by $\widehat{\pi}_R$, $\widehat{\pi}_B$ and \widehat{H}_B (and P).

Now let Γ be the graph with vertex set $\binom{[m]}{2}$ and edge set Π_Γ . Then $e(\Gamma) = |\Pi_\Gamma| \geq \gamma n^2$. We furthermore claim that $\Delta_\Gamma \leq 4n^2 / m^2$. Indeed, for any given $\{x, y\} \in \binom{[m]}{2}$, due to \mathcal{F}_R holding for $\pi_R = \widehat{\pi}_R$, we have $|(\widehat{\pi}_R)^{-1}(x)| \leq 2n/m$ and $|(\widehat{\pi}_R)^{-1}(y)| \leq 2n/m$, and hence there can be at most $(2n/m)^2 = 4n^2 / m^2$ pairs $\{i, j\} \subseteq [s]$ with $\{\widehat{\pi}_R(u_i), \widehat{\pi}_R(u_j)\} = \{x, y\}$ or $\{\widehat{\pi}_R(v_i), \widehat{\pi}_R(v_j)\} = \{x, y\}$. Thus, $\{x, y\}$ appears in at most $4n^2 / m^2$ of the pairs $\{\{\widehat{\pi}_R(u_i), \widehat{\pi}_R(u_j)\}, \{\widehat{\pi}_R(v_i), \widehat{\pi}_R(v_j)\}\}$ forming Π . In particular, each $\{x, y\} \in \binom{[m]}{2}$ appears in at most $4n^2 / m^2$ pairs in Π_Γ , and hence $\{x, y\}$ as a vertex of Γ has degree at most $4n^2 / m^2$. This shows $\Delta_\Gamma \leq 4n^2 / m^2$, as claimed.

Thus, we can apply Lemma 4.4 to the graph Γ on the vertex set $\binom{[m]}{2}$. We find that

$$\mathbb{P}[\mathcal{T}_R \text{ holds and there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H'_R)] \leq n^{-2n}. \quad (2)$$

Finally, recall that our goal was to show that, conditional on our fixed outcomes of π_R , π_B and H_B we have (see (1))

$$\mathbb{P}[\mathcal{F}_R, \mathcal{F}_B, \mathcal{D}_R, \mathcal{D}_B, \mathcal{T}_R, \mathcal{T}_B, \mathcal{E}(\varepsilon)] \text{ hold}$$

$$\text{and } P \text{ is an odd connected pairing in } G \mid \pi_R = \widehat{\pi}_R, \pi_B = \widehat{\pi}_B, H_B = \widehat{H}_B \leq n^{-2n}.$$

We claim that whenever there is a pair $\{e, f\} \in E(\Gamma)$ with $e, f \in E(H'_R)$, and $\pi_R = \widehat{\pi}_R$, $\pi_B = \widehat{\pi}_B$, and $H_B = \widehat{H}_B$, the pairing P cannot be an odd connected pairing in G . Indeed, as $\{e, f\} \in E(\Gamma) = \Pi_\Gamma$, we must have $\{e, f\} = \{\{\pi_R(u_i), \pi_R(u_j)\}, \{\pi_R(v_i), \pi_R(v_j)\}\}$ for some indices $i, j \in [s]$ with $\pi_R(u_i) \neq \pi_R(u_j)$ and $\pi_R(v_i) \neq \pi_R(v_j)$ and $\{\pi_R(u_i), \pi_R(u_j)\} \neq \{\pi_R(v_i), \pi_R(v_j)\}$, such that neither $\{u_i, u_j\}$ nor $\{v_i, v_j\}$ is closed by a blue cherry in H_0 . As $e, f \in E(H'_R)$, we now have $\{\pi_R(u_i), \pi_R(u_j)\}, \{\pi_R(v_i), \pi_R(v_j)\} \in E(H'_R)$. By construction of H_0 , this means that $\{u_i, u_j\}, \{v_i, v_j\} \in E(H_0)$ are red edges of H_0 . Since neither $\{u_i, u_j\}$ nor $\{v_i, v_j\}$ is closed by a blue cherry, neither of these red edges gets deleted when passing from H_0 to H and so we have $\{u_i, u_j\}, \{v_i, v_j\} \in E(H)$. Recalling that G is the complement of the graph H , this means that $\{u_i, u_j\}, \{v_i, v_j\} \notin E(G)$. Therefore P is not an odd connected pairing in G . Thus, we can conclude that

$$\mathbb{P}[\mathcal{F}_R, \mathcal{F}_B, \mathcal{D}_R, \mathcal{D}_B, \mathcal{T}_R, \mathcal{T}_B, \mathcal{E}(\varepsilon)] \text{ hold}$$

$$\text{and } P \text{ is an odd connected pairing in } G \mid \pi_R = \widehat{\pi}_R, \pi_B = \widehat{\pi}_B, H_B = \widehat{H}_B$$

$$\leq \mathbb{P}[\mathcal{T}_R \text{ holds and } P \text{ is an odd connected pairing in } G \mid \pi_R = \widehat{\pi}_R, \pi_B = \widehat{\pi}_B, H_B = \widehat{H}_B]$$

$$\leq \mathbb{P}[\mathcal{T}_R \text{ holds and there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H'_R) \mid \pi_R = \widehat{\pi}_R, \pi_B = \widehat{\pi}_B, H_B = \widehat{H}_B]$$

$$= \mathbb{P}[\mathcal{T}_R \text{ holds and there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H'_R)]$$

$$\leq n^{-2n},$$

where in the penultimate step we used that the events \mathcal{T}_R and $e, f \in E(H'_R)$ for $\{e, f\} \in E(\Gamma)$ only depend on the outcome of the random graphs H_R^* and H_R , and are therefore not affected by the conditioning on the outcomes of π_R, π_B, H_B (which are independent of H_R^* and H_R). Furthermore, in the last step we used (2). This is precisely what we wanted to prove. \square

5.3. Proof of Theorems 1.3 and 1.4. Finally, in order to complete the proofs of Theorems 1.3 and 1.4, in this section we give the short deterministic argument showing that Theorem 3.1 implies Theorem 1.4. At the end of this section, we also include the almost-immediate deduction of Theorem 1.3 from Theorem 1.4.

Proof of Theorem 1.4, using Theorem 3.1. Fix $0 < \eta < 1$. Our goal is to show that if n is sufficiently large, there is an n -vertex graph G with $\alpha(G) \leq 2$ which does not have K_t as an odd minor for $t = \lceil (\frac{1}{3} + \eta)n \rceil$. Let $\varepsilon = \eta/2$, and let G be an n -vertex graph given by Theorem 3.1, so that $\alpha(G) \leq 2$ and G does not have an odd connected pairing of size at least εn . It remains to prove that G does not have K_t as an odd minor.

So suppose for contradiction that there are vertex-disjoint trees $T_1, \dots, T_t \subseteq G$ and colorings χ_1, \dots, χ_t as in the definition of an odd minor. Then each of the trees T_1, \dots, T_t has at least one vertex. Furthermore $|V(T_1)| + \dots + |V(T_t)| \leq n$, and hence there can be at most $n/3$ indices $i \in [t]$ with $|V(T_i)| \geq 3$. As a consequence, at least $t - (n/3) \geq (\frac{1}{3} + \eta)n - (n/3) = \eta n$ of the trees T_1, \dots, T_t have at most two vertices. In particular, letting $s = \lceil \eta n / 2 \rceil < n/2$, we find that at least s of the trees have exactly two vertices, or at least s of the trees have exactly one vertex.

Suppose that the first happens, so without loss of generality T_1, \dots, T_s have exactly two vertices each. For every $i \in [s]$, we know that χ_i is a proper 2-coloring of the edge T_i , hence one vertex of T_i receives color white and the other color black. We can thus define u_i to be the

white vertex of T_i , and v_i the black vertex. Then $\{(u_1, v_1), \dots, (u_s, v_s)\}$ define an odd connected pairing of size s : for all $i \neq j$, we know that there is some monochromatic (with respect to χ_i, χ_j) edge between T_i and T_j , and this edge must be either $\{u_i, u_j\}$ or $\{v_i, v_j\}$.

In the second case, suppose without loss of generality that T_1, \dots, T_s each consist of a single vertex. Then the fact that these s vertices form an odd K_s minor implies that, for every $i \neq j$, there is an edge between the single vertex of T_i and the single vertex of T_j . In particular, we have actually found that $K_s \subseteq G$. Now, label the vertices in T_1, \dots, T_s as u_1, \dots, u_s , and choose s other arbitrary vertices v_1, \dots, v_s , which we may do since $s < n/2$. Then $\{(u_1, v_1), \dots, (u_s, v_s)\}$ is an odd connected pairing since for any distinct $i, j \in [s]$ there is an edge between u_i and u_j .

Thus, in either case, we have found that G contains an odd connected pairing of size s . Noting that $s \geq \eta n/2 = \varepsilon n$, this is a contradiction to our choice of G . \square

We end this section by showing that Theorem 1.3 follows from Theorem 1.4.

Proof of Theorem 1.3, using Theorem 1.4. Let $\delta > 0$, and without loss of generality assume $\delta < 1$. Set $\eta := \delta/5 > 0$. Now, for any sufficiently large t , we can apply Theorem 1.4 to $n = \lceil (3 - 10\eta)t \rceil$ and η , and obtain an n -vertex graph G with $\alpha(G) \leq 2$ not having K_t as an odd minor (here, we used that $t \geq (\frac{1}{3} + \eta)(3 - 9\eta)t \geq (\frac{1}{3} + \eta)n$, if t is sufficiently large). We conclude by observing that $\chi(G) \geq n/\alpha(G) \geq n/2 \geq (\frac{3}{2} - 5\eta)t = (\frac{3}{2} - \delta)t$. \square

6. PROPERTIES OF THE RANDOM FUNCTIONS

In this section, we provide the proof of Lemma 4.2.

Proof of Lemma 4.2. Fix $0 < \varepsilon < 1/10$, and let $s = \lceil \varepsilon n \rceil$. We need to show that as $n \rightarrow \infty$, it holds with high probability that for each s -pairing P in $[n]$ at least one of the random functions π_R and π_B is ε -respectful of P . We may in particular assume that n is sufficiently large with respect to ε .

We begin by fixing an s -pairing $P = \{(u_1, v_1), \dots, (u_s, v_s)\}$ in $[n]$, and letting $\pi : [n] \rightarrow [m]$ be a uniformly random function. We claim that

$$\mathbb{P}[\pi \text{ is not } \varepsilon\text{-respectful of } P] \leq n^{-(5/8)\varepsilon n}. \quad (3)$$

First, note that if π is not ε -respectful of P , then this means that there is a set $T \subseteq \binom{[m]}{2}$ of size $|T| = \lfloor \varepsilon s \rfloor$ such that for all $i \in [s]$ we have $\{\pi(u_i), \pi(v_i)\} \in T$ or $\pi(u_i) = \pi(v_i)$.

For a fixed set $T \subseteq \binom{[m]}{2}$ of size $|T| = \lfloor \varepsilon s \rfloor$ and a fixed index $i \in [s]$, we have

$$\mathbb{P}[\{\pi(u_i), \pi(v_i)\} \in T \text{ or } \pi(u_i) = \pi(v_i)] \leq \mathbb{P}[\pi(u_i) = \pi(v_i)] + \mathbb{P}[\{\pi(u_i), \pi(v_i)\} \in T].$$

The first summand on the right-hand side is exactly $1/m$, and the second summand is exactly $2|T|/m^2$. Indeed, $(\pi(u_i), \pi(v_i)) \in [m]^2$ is uniformly distributed among all m^2 options, and precisely $2|T|$ of these options correspond to pairs in $T \subseteq \binom{[m]}{2}$ (note that for every pair in T there are two possible orderings, meaning that there are two corresponding elements of $[m]^2$). In total, we find that

$$\mathbb{P}[\{\pi(u_i), \pi(v_i)\} \in T \text{ or } \pi(u_i) = \pi(v_i)] \leq \frac{1}{m} + \frac{2|T|}{m^2} = \frac{2|T| + m}{m^2} \leq \frac{2\varepsilon^2 n + 2 + m}{m^2} \leq \frac{3\varepsilon^2 n}{m^2} \leq n^{-7/8}$$

for every $i \in [s]$ (here, we used that $|T| = \lfloor \varepsilon s \rfloor \leq \varepsilon(\varepsilon n + 1) < \varepsilon^2 n + 1$ and $m = \lceil n/(\log n)^8 \rceil$, and that n is sufficiently large with respect to ε). Now, the outcomes of $\pi(u_i)$ and $\pi(v_i)$ are independent over all the different $i \in [s]$. Therefore, we obtain

$$\mathbb{P}[\text{for all } i \in [s] \text{ we have } \{\pi(u_i), \pi(v_i)\} \in T \text{ or } \pi(u_i) = \pi(v_i)] \leq (n^{-7/8})^s \leq n^{-(7/8)\varepsilon n}.$$

Applying the union bound over the

$$\binom{\binom{m}{2}}{\lfloor \varepsilon s \rfloor} \leq \binom{m}{2}^{\varepsilon s} \leq (n^2)^{\varepsilon^2 n+1} = n^{2\varepsilon^2 n+2} \leq n^{\varepsilon n/4}$$

choices of $T \subseteq \binom{[m]}{2}$ of size $|T| = \lfloor \varepsilon s \rfloor$ (where we again used that n is sufficiently large with respect to ε), we find that

$$\mathbb{P}[\pi \text{ is not } \varepsilon\text{-respectful of } P] \leq n^{\varepsilon n/4} \cdot n^{-(7/8)\varepsilon n} = n^{-(5/8)\varepsilon n},$$

proving (3). As a consequence of (3), we see that for the fixed pairing P , we have

$$\mathbb{P}[\pi_R \text{ and } \pi_B \text{ are both not } \varepsilon\text{-respectful of } P] = (\mathbb{P}[\pi \text{ is not } \varepsilon\text{-respectful of } P])^2 \leq n^{-(5/4)\varepsilon n},$$

since π_R, π_B are independent and both uniformly random among all maps $[n] \rightarrow [m]$. Finally, we note that the number of s -pairings P in $[n]$ is at most $\binom{n}{s}^2 s!$, as we may pick the set $\{u_1, \dots, u_s\}$, then pick the set $\{v_1, \dots, v_s\}$, and finally pick a bijection between them encoding the pairs $\{u_i, v_i\}$. This number is

$$\binom{n}{s}^2 s! \leq 2^{2n} s^s \leq 2^{2n} n^s \leq 2^{2n} n^{\varepsilon n+1} = n^{\varepsilon n+1+2n/\log_2 n} \leq n^{(9/8)\varepsilon n}.$$

Applying a union bound over all choices of P , we find that

$$\begin{aligned} \mathbb{P}[\text{for some } s\text{-pairing } P \text{ in } [n] \text{ neither } \pi_R \text{ nor } \pi_B \text{ is } \varepsilon\text{-respectful of } P] \\ \leq n^{(9/8)\varepsilon n} \cdot n^{-(5/4)\varepsilon n} = n^{-\varepsilon n/8}. \end{aligned}$$

Thus, for n sufficiently large with respect to ε , the probability that the event $\mathcal{E}(\varepsilon)$ fails is at most $n^{-\varepsilon n/8}$. In particular, $\mathcal{E}(\varepsilon)$ holds with high probability as $n \rightarrow \infty$. \square

7. PROPERTIES OF THE RANDOM GRAPHS

Our main goal in this section is to prove Lemma 4.4; in Section 7.3 we also prove the simple Lemmas 4.1 and 4.3.

7.1. A preparatory lemma. In order to prove Lemma 4.4, we will need Lemma 7.2 below. Before stating and proving this lemma, we state a concentration inequality, closely related to Talagrand's inequality [52], which was proved by McDiarmid and Reed [37, Theorem 2].

Lemma 7.1. *Let $r > 0$, let X_1, \dots, X_N be independent random variables taking values in a finite set S , and let $X = f(X_1, \dots, X_N)$ for some $f: S^N \rightarrow \mathbb{R}_{\geq 0}$. Suppose that for all $i \in [N]$ and $\mathbf{x} \in S^N$, changing the i -th coordinate of \mathbf{x} changes the value of $f(\mathbf{x})$ by at most 1. Furthermore, suppose that for every $\mathbf{x} \in S^N$ there exists a subset $J \subseteq [N]$ with $|J| \leq r f(\mathbf{x})$ such that every $\mathbf{y} \in S^N$ with $\mathbf{y}_i = \mathbf{x}_i$ for all $i \in J$ satisfies $f(\mathbf{y}) \geq f(\mathbf{x})$. Then, for all $t \geq 0$, we have*

$$\mathbb{P}[X \leq \mathbb{E}[X] - t] \leq \exp\left(-\frac{t^2}{2 \cdot (r\mathbb{E}[X] + t/3)}\right).$$

The key lemma in the proof of Lemma 4.4 is the following lemma about a random vertex subset V^* of a given graph Γ , containing each vertex of Γ independently with some probability q . The lemma states that, with large probability, in $\Gamma[V^*]$ we can find a large matching (i.e. a large set of pairwise disjoint edges). We will apply this lemma with Γ being the graph on $\binom{[m]}{2}$ given in Lemma 4.4, and the random subset $V^* \subseteq \binom{[m]}{2}$ corresponding to the edge set of the random graph H_X^* .

Lemma 7.2. *Let Γ be a graph, and consider a random vertex set $V^* \subseteq V(\Gamma)$ where every vertex of Γ is included independently with some probability $q \leq 1/\Delta_\Gamma$. Then with probability at least*

$$1 - \exp\left(-\frac{q^2 e(\Gamma)}{200}\right),$$

there exists a matching M^ of size at least $q^2 e(\Gamma)/20$ in the induced subgraph $\Gamma[V^*]$.*

To prove the lemma, we apply Lemma 7.1 to the matching number of $\Gamma[V^*]$.

Proof of Lemma 7.2. If $e(\Gamma) = 0$, the statement in the lemma trivially holds, so we may assume $e(\Gamma) \geq 1$, and hence $\Delta_\Gamma \geq 1$.

We need to show that with very large probability the graph $\Gamma[V^*]$ contains a matching of size at least $q^2 e(\Gamma)/20$, i.e., it has matching number $\nu(\Gamma[V^*]) \geq q^2 e(\Gamma)/20$. To do so, we first show a lower bound on $\mathbb{E}[\nu(\Gamma[V^*])]$. Let us consider the set E^* consisting of all edges $\{u, v\} \in E(\Gamma[V^*])$ with $d_{\Gamma[V^*]}(u) = d_{\Gamma[V^*]}(v) = 1$. In other words, E^* is the collection of all edges $\{u, v\} \in E(\Gamma)$ with $u, v \in V^*$ where both u and v have degree 1 in $\Gamma[V^*]$. Note that then the edges $e \in E^*$ are automatically pairwise disjoint (because if two such edges were to meet in a vertex, then this vertex would need to have degree at least 2), and thus we always have $\nu(\Gamma[V^*]) \geq |E^*|$. Observe that for any fixed edge $\{u, v\} \in E(\Gamma)$ we have

$$\begin{aligned} \mathbb{P}[\{u, v\} \in E^*] &= \mathbb{P}[u, v \in V^* \text{ and } d_{\Gamma[V^*]}(u) = d_{\Gamma[V^*]}(v) = 1] \\ &\geq \mathbb{P}[u, v \in V^* \text{ and } x \notin V^* \text{ for all } x \in N_\Gamma(u) \cup N_\Gamma(v) \setminus \{u, v\}] \\ &\geq q^2 \cdot (1 - q)^{2(\Delta_\Gamma - 1)} \geq q^2 \cdot \left(1 - \frac{1}{\Delta_\Gamma}\right)^{2(\Delta_\Gamma - 1)} \geq q^2 \cdot e^{-2} \geq \frac{q^2}{10}, \end{aligned}$$

where we used the assumption $q \leq 1/\Delta_\Gamma$ and that $(1 - 1/t)^{t-1} \geq 1/e$ holds for all positive integers t . Thus, we obtain

$$\mathbb{E}[\nu(\Gamma[V^*])] \geq \mathbb{E}[|E^*|] \geq e(\Gamma) \cdot \frac{q^2}{10} = \frac{q^2 e(\Gamma)}{10}.$$

Note that we can view $\nu(\Gamma[V^*])$ as a function of the indicator random variables for the events $v \in V^*$ for $v \in V(\Gamma)$. Further note that changing only one of these variables can change $\nu(\Gamma[V^*])$ by at most 1 (since adding or omitting a vertex from V^* changes the matching number of $\Gamma[V^*]$ by at most 1). We claim that, taking $r = 2$, also the second condition in Lemma 7.1 is satisfied. Indeed, for every vertex subset $X \subseteq V(\Gamma)$, we can take $J \subseteq X \subseteq V(\Gamma)$ to be the vertex set of a matching of size $\nu(\Gamma[X])$ in $\Gamma[X]$. Then we have $|J| = 2\nu(\Gamma[X])$, and we furthermore have $\nu(\Gamma[Y]) \geq \nu(\Gamma[X])$ for every $Y \subseteq V(\Gamma)$ with $J \subseteq Y$. Hence, we can apply Lemma 7.1 to the random variable $\nu(\Gamma[V^*])$ with the parameters $r = 2$ and $t = \mathbb{E}[\nu(\Gamma[V^*])]/2$ to obtain:

$$\begin{aligned} \mathbb{P}\left[\nu(\Gamma[V^*]) < \frac{q^2 e(\Gamma)}{20}\right] &\leq \mathbb{P}\left[\nu(\Gamma[V^*]) < \mathbb{E}[\nu(\Gamma[V^*])] - \frac{\mathbb{E}[\nu(\Gamma[V^*])]}{2}\right] \\ &\leq \exp\left(-\frac{(\mathbb{E}[\nu(\Gamma[V^*])]/2)^2}{2 \cdot (2\mathbb{E}[\nu(\Gamma[V^*])] + \mathbb{E}[\nu(\Gamma[V^*])]/6)}\right) \\ &\leq \exp\left(-\frac{\mathbb{E}[\nu(\Gamma[V^*])]}{20}\right) \leq \exp\left(-\frac{q^2 e(\Gamma)}{200}\right). \quad \square \end{aligned}$$

7.2. Proof of Lemma 4.4. We are now ready to prove Lemma 4.4. In the proof, we apply Lemma 7.2 to the given graph Γ with vertex set $\binom{[m]}{2}$ and the random graph H_X^* , whose edges can be viewed as a random subset of the vertex set $V(\Gamma)$. We can then conclude that with very large probability there exists a large matching M^* in Γ such that $e, f \in E(H_X^*)$ for all $\{e, f\} \in M$. Therefore, we may condition on a fixed outcome of H_X^* such that this holds, and

such that the event \mathcal{T}_X holds (meaning that every edge in H_X^* is contained in only few triangles). We can then select a submatching $M \subseteq M^*$ such that not just all $\{e, f\} \in M^*$ are disjoint from one another, but also the sets of those edges that together with e or f lie on a triangle in H_X^* are disjoint for all $\{e, f\} \in M$. For such a submatching $M \subseteq M^*$, the events of having $e, f \in E(H_X')$ (i.e., having $e, f \in E(H_X)$ and neither e nor f lying on a triangle in H_X) are independent for all $\{e, f\} \in M$. This allows us to conclude that at least one of these events is likely to happen.

Proof of Lemma 4.4. Without loss of generality, let $X = R$ (the case $X = B$ is analogous). We are given a graph Γ with vertex set $\binom{[m]}{2}$ such that $e(\Gamma) \geq \gamma n^2$ and $\Delta_\Gamma \leq 4n^2/m^2 \leq 4 \log^{16} n \leq 8 \log^{16} m$. The edge set of the random graph $E(H_R^*)$ can be interpreted as a random subset of $\binom{[m]}{2}$ which includes every element of $\binom{[m]}{2}$ independently with probability $m^{-1/2}$. Thus, $E(H_R^*)$ is a random subset of the vertex set of Γ , and we can apply Lemma 7.2, noting that if n (and thus m) is sufficiently large, we have $m^{-1/2} \leq 1/(8 \log^{16} m) \leq 1/\Delta_\Gamma$. Let \mathcal{H} denote the event (depending only on the outcome of H_R^*) that there exists a matching M^* of size $|M^*| \geq (m^{-1/2})^2 e(\Gamma)/20 = e(\Gamma)/(20m)$ in the induced subgraph $\Gamma[E(H_R^*)]$, i.e. a matching M^* of size $|M^*| \geq e(\Gamma)/(20m)$ in Γ such that $e, f \in E(H_R^*)$ for all $\{e, f\} \in M^*$. By Lemma 7.2 we have

$$\mathbb{P}[\mathcal{H}] \geq 1 - \exp\left(-\frac{(m^{-1/2})^2 e(\Gamma)}{200}\right) \geq 1 - \exp\left(-\frac{\gamma n^2}{200m}\right) \geq 1 - \exp(-n \log^7 n),$$

recalling that $m = \lceil n/\log^8 n \rceil$ and n is sufficiently large with respect to γ .

It now suffices to show that

$$\mathbb{P}[\text{there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H_R') \mid \mathcal{H} \text{ and } \mathcal{T}_R] \leq \exp(-3n \log n). \quad (4)$$

Indeed, upon proving this we can conclude that

$$\begin{aligned} & \mathbb{P}[\mathcal{T}_R \text{ holds and there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H_R')] \\ & \leq \mathbb{P}[\mathcal{H} \text{ and } \mathcal{T}_R \text{ hold and there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H_R')] + \mathbb{P}[\mathcal{H}^c] \\ & \leq \mathbb{P}[\text{there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H_R') \mid \mathcal{H} \text{ and } \mathcal{T}_R] + \mathbb{P}[\mathcal{H}^c] \\ & \leq \exp(-3n \log n) + \exp(-n \log^7 n) \leq \exp(-2n \log n) = n^{-2n}, \end{aligned}$$

as desired. So it only remains to prove (4).

Recall that both of the events \mathcal{H} and \mathcal{T}_R are determined by the outcome of the random graph H_R^* . Thus, to show (4), we can condition on a fixed outcome \widehat{H}_R^* of H_R^* satisfying \mathcal{H} and \mathcal{T}_R . This means that there is a matching M^* in Γ of size $|M^*| \geq e(\Gamma)/(20m) \geq \gamma n^2/(20m)$ with $e, f \in E(\widehat{H}_R^*)$ for all $\{e, f\} \in M^*$, and furthermore that every edge in \widehat{H}_R^* lies on at most $20 \log m$ triangles of \widehat{H}_R^* . In particular, for any edge $e \in E(\widehat{H}_R^*)$ there are at most $40 \log m$ other edges appearing on a triangle in \widehat{H}_R^* together with e . For any $\{e, f\} \in M^*$, let $S_{\{e, f\}} \subseteq E(\widehat{H}_R^*)$ be the set of edges that are equal to e or f or appear on a triangle in \widehat{H}_R^* together with e or f . Then we have $|S_{\{e, f\}}| \leq 40 \log m + 40 \log m + 2 \leq 100 \log m$ for any $\{e, f\} \in M^*$, and furthermore every edge appears in $S_{\{e, f\}}$ for at most $40 \log m + 1 \leq 50 \log m$ different $\{e, f\} \in M^*$.

Now, for any $\{e, f\} \in M^*$ there are at most $100 \log m \cdot 50 \log m = 5000 \log^2 m$ different $\{e', f'\} \in M^*$ with $S_{\{e, f\}} \cap S_{\{e', f'\}} \neq \emptyset$. Therefore, via a simple greedy procedure (or alternatively by the bound of the independence number of a graph in terms of its maximum degree), we can choose a submatching $M \subseteq M^*$ of size

$$|M| \geq \frac{|M^*|}{5000 \log^2 m + 1} \geq \frac{\gamma n^2/(20m)}{10^4 \log^2 m} \geq \frac{\gamma n^2}{10^6 m \log^2 m} \geq \frac{\gamma}{2 \cdot 10^6} \cdot n \log^6 n \geq 6n \log^5 n$$

such that for any distinct $\{e, f\}, \{e', f'\} \in M$ we have $S_{\{e, f\}} \cap S_{\{e', f'\}} = \emptyset$.

Now, conditional on the event $H_R^* = \widehat{H_R^*}$, the random graph H_R is obtained as a random subgraph $H_R \subseteq \widehat{H_R^*}$ that includes every edge of $\widehat{H_R^*}$ independently with probability $1/\log^2 m$. Note that for every $\{e, f\} \in M \subseteq M^* \subseteq E(\Gamma)$, whenever we have $S_{\{e, f\}} \cap E(H_R) = \{e, f\}$, we have $e, f \in E(H_R)$ and no triangle in H_R contains e or f . Now, the events $S_{\{e, f\}} \cap E(H_R) = \{e, f\}$ are independent for all $\{e, f\} \in M$, and for each $\{e, f\} \in M$ we have

$$\begin{aligned} \mathbb{P}[S_{\{e, f\}} \cap E(H_R) = \{e, f\} \mid H_R^* = \widehat{H_R^*}] &\geq \left(\frac{1}{\log^2 m}\right)^2 \cdot \left(1 - \frac{1}{\log^2 m}\right)^{|S_{\{e, f\}}|-2} \\ &\geq \frac{1}{\log^4 m} \cdot \left(1 - \frac{1}{\log^2 m}\right)^{100 \log m} \geq \frac{1}{2 \log^4 m}. \end{aligned}$$

Thus, we obtain

$$\begin{aligned} &\mathbb{P}[\text{there is no } \{e, f\} \in M \text{ with } e, f \in E(H_R) \text{ s.t. no triangle in } H_R \text{ contains } e \text{ or } f \mid H_R^* = \widehat{H_R^*}] \\ &\leq \mathbb{P}[\text{there is no } \{e, f\} \in M \text{ with } S_{\{e, f\}} \cap E(H_R) = \{e, f\} \mid H_R^* = \widehat{H_R^*}] \\ &\leq \left(1 - \frac{1}{2 \log^4 m}\right)^{|M|} \leq \exp\left(-\frac{|M|}{2 \log^4 m}\right) \leq \exp\left(-\frac{6n \log^5 n}{2 \log^4 m}\right) \leq \exp(-3n \log n). \end{aligned}$$

For any $\{e, f\} \in M \subseteq M^* \subseteq E(\Gamma)$ with $e, f \in E(H_R)$ such that no triangle in H_R contains e or f , by the definition of H'_R we have $e, f \in E(H'_R)$. Thus, we can conclude

$$\mathbb{P}[\text{there is no } \{e, f\} \in E(\Gamma) \text{ with } e, f \in E(H'_R) \mid H_R^* = \widehat{H_R^*}] \leq \exp(-3n \log n),$$

as desired. \square

7.3. Proof of Lemmas 4.1 and 4.3. In this subsection we derive Lemmas 4.1 and 4.3 from the Chernoff bound. We will use the following simple form of the Chernoff bound, which follows for example from [25, Theorem 2.8 and eq. (2.10)].

Lemma 7.3. *Suppose that X is a sum of independent indicator random variables. Then, for all $t \geq 2 \cdot \mathbb{E}[X]$, we have*

$$\mathbb{P}[X \geq t] \leq \exp\left(-\frac{t}{6}\right).$$

With this tool, the proofs of Lemmas 4.1 and 4.3 are quite straightforward.

Proof of Lemma 4.1. It suffices to prove the results for \mathcal{D}_R and \mathcal{F}_R , since the proofs for \mathcal{D}_B and \mathcal{F}_B are analogous. First, fix some vertex $v \in V_R$, and let X_v be its degree in H_R . Then X_v is distributed as a binomial random variable with parameters $(m-1, p)$, so in particular $\mathbb{E}[X_v] \leq pm = (2pm)/2$. Thus, Lemma 7.3 implies $\mathbb{P}[X_v \geq 2pm] \leq \exp(-2pm/6) = \exp(-m^{1/2} \log^{-2} m/3) \leq m^{-2}$, where the final inequality holds for n and hence m sufficiently large. Thus, taking a union bound over all $v \in V_R$ proves that \mathcal{D}_R holds with high probability.

Similarly, for fixed $v \in V_R$, let Y_v denote the random variable $|\pi_R^{-1}(v)|$. We thus have that Y_v is distributed as a binomial random variable with parameters $(n, 1/m)$, so $\mathbb{E}[Y_v] = n/m$. We may thus apply Lemma 7.3 to see that $\mathbb{P}[Y_v \geq 2n/m] \leq \exp(-(2n/m)/6) \leq \exp(-\log^8 n/4) \leq n^{-2} \leq m^{-2}$, where the penultimate inequality again holds for sufficiently large n . Another union bound over all $v \in V_R$ shows that \mathcal{F}_R holds with high probability. \square

Proof of Lemma 4.3. It suffices to show that \mathcal{T}_R holds with high probability; the proof for \mathcal{T}_B is analogous. For distinct $u, v \in V(H_R^*)$, the number $X_{u,v}$ of common neighbors of u and v in H_R^* is a binomial random variable with parameters $(m-2, m^{-1})$. Indeed, any vertex $w \in V(H_R^*) \setminus \{u, v\}$ has edges to both u and v with probability $(m^{-1/2})^2 = m^{-1}$, and this

happens independently for all $w \in V(H_R^*) \setminus \{u, v\}$. Since $\mathbb{E}[X_{u,v}] \leq 1 \leq (20 \log m)/2$, Lemma 7.3 yields $\mathbb{P}[X_{u,v} \geq 20 \log m] \leq \exp(-3 \log m) = m^{-3}$. Taking a union bound over the at most m^2 choices for $u, v \in V(H_R^*)$ shows that with high probability we have $X_{u,v} \leq 20 \log m$ for all distinct $u, v \in V(H_R^*)$. Whenever this happens, every edge in H_R^* is contained in at most $20 \log m$ triangles, and so \mathcal{T}_R holds. \square

8. CONCLUDING REMARKS

Given that we have disproved the odd Hadwiger conjecture, and given the close connection between Hadwiger's conjecture and its odd variant, it is natural to wonder whether our approach could disprove Hadwiger's conjecture itself. At the moment, there appear to be serious obstructions to achieving this.

The most natural thing one could try in order to disprove Hadwiger's conjecture is to simply analyze the presence of $K_{n/2}$ as a minor in our construction, and aim to prove that with high probability, the constructed graph does not have $K_{n/2}$ as a minor. There are at least two serious issues with making such an approach work. The first concerns the numbers: if we now consider $K_{n/2}$ as an ordinary minor (no longer required to be odd), then we have four potential connections available between two given trees T_i and T_j of order 2. We would thus want $p^4 n^2 \gg n \log n$ for our union bound to work, requiring $p \gg n^{-1/4}$. Since we also need $p \leq c/\sqrt{m}$ in order to delete all triangles from $G(m, p)$ without removing many edges, this would require $m \ll \sqrt{n}$; however, several parts of our proof completely break down at such a small value of m .

The other issue seems perhaps even more serious. Recall that in our proof, we immediately dropped one of the conditions defining an odd connected matching, and instead passed to working with an odd connected pairing; namely, we did not require that our branch sets T_i actually define a connected subgraph of G . For odd minors, dropping this condition turned out to be immaterial. However, for ordinary minors, one cannot simply drop this condition: it is easy to see that in any proper coloring of G with $\chi(G)$ colors, there must be at least one edge between every two color classes. Thus, if we allow branch sets that are not connected, then every graph G contains such a "weak minor" of order $\chi(G)$, simply obtained by contracting each color class of any optimal coloring. Thus, any argument that aims to disprove Hadwiger's conjecture must crucially use that the branch sets are connected, and it is difficult for us to see how this extra condition could significantly affect the relevant probabilities.

Acknowledgments: We would like to thank Micha Christoph, Jacob Fox, Marc Kaufmann, Anders Martinsson, and Rob Morris for helpful and inspiring discussions related to this topic.

REFERENCES

- [1] M. Ajtai, J. Komlós, and E. Szemerédi, A note on Ramsey numbers, *J. Combin. Theory Ser. A* **29** (1980), 354–360.
- [2] K. Appel and W. Haken, Every planar map is four colorable. I. Discharging, *Illinois J. Math.* **21** (1977), 429–490.
- [3] K. Appel, W. Haken, and J. Koch, Every planar map is four colorable. II. Reducibility, *Illinois J. Math.* **21** (1977), 491–567.
- [4] T. Bohman, The triangle-free process, *Adv. Math.* **221** (2009), 1653–1677.
- [5] T. Bohman and P. Keevash, Dynamic concentration of the triangle-free process, *Random Structures Algorithms* **58** (2021), 221–293.
- [6] B. Bollobás, *Random graphs, Cambridge Studies in Advanced Mathematics*, vol. 73, second ed., Cambridge University Press, Cambridge, 2001.

- [7] M. Campos, M. Jenssen, M. Michelen, F. Pfender, and J. Sahasrabudhe, A polynomial improvement for the odd cycle-complete Ramsey numbers, 2025. Preprint available at arXiv:2511.10641.
- [8] M. Campos, M. Jenssen, M. Michelen, and J. Sahasrabudhe, A new lower bound for the Ramsey numbers $R(3, k)$, 2025. Preprint available at arXiv:2505.13371.
- [9] P. A. Catlin, Hajós' graph-coloring conjecture: variations and counterexamples, *J. Combin. Theory Ser. B* **26** (1979), 268–274.
- [10] M. Chudnovsky, Hadwiger's conjecture and seagull packing, *Notices Amer. Math. Soc.* **57** (2010), 733–736.
- [11] Y. Colin de Verdière, Sur un nouvel invariant des graphes et un critère de planarité, *J. Combin. Theory Ser. B* **50** (1990), 11–21.
- [12] J. Costa, E. Luu, D. R. Wood, and J. H. Yip, Verifying Hadwiger's conjecture for examples of graphs with $\alpha(G) = 2$, 2025. Preprint available at arXiv:2512.17114.
- [13] M. Delcourt and L. Postle, Reducing linear Hadwiger's conjecture to coloring small graphs, *J. Amer. Math. Soc.* **38** (2025), 481–507.
- [14] P. Duchet and H. Meyniel, On Hadwiger's number and the stability number, in *Graph theory (Cambridge, 1981)*, *North-Holland Math. Stud.*, vol. 62, North-Holland, Amsterdam-New York, 1982, 71–73.
- [15] P. Erdős, Graph theory and probability. II, *Canadian J. Math.* **13** (1961), 346–352.
- [16] P. Erdős, S. Suen, and P. Winkler, On the size of a random maximal graph, *Random Structures Algorithms* **6** (1995), 309–318.
- [17] G. Fiz Pontiveros, S. Griffiths, and R. Morris, The triangle-free process and the Ramsey number $R(3, k)$, *Mem. Amer. Math. Soc.* **263** (2020), v+125.
- [18] J. Fox, Complete minors and independence number, *SIAM J. Discrete Math.* **24** (2010), 1313–1321.
- [19] T. Gallai, Kritische Graphen. II, *Magyar Tud. Akad. Mat. Kutató Int. Közl.* **8** (1963), 373–395.
- [20] J. Geelen, B. Gerards, B. Reed, P. Seymour, and A. Vetta, On the odd-minor variant of Hadwiger's conjecture, *J. Combin. Theory Ser. B* **99** (2009), 20–29.
- [21] B. Guenin, Odd- K_5 -free graphs are 4-colourable, in R. Diestel, A. Schrijver, and P. Seymour (eds.), *Oberwolfach Reports*, vol. 2, EMS Press, Berlin, 2005, 176–178.
- [22] H. Hadwiger, Über eine Klassifikation der Streckenkomplexe, *Vierteljahrsschr. Naturforsch. Ges. Zürich* **88** (1943), 133–142.
- [23] Z. Hefty, P. Horn, D. King, and F. Pfender, Improving $R(3, k)$ in just two bites, 2025. Preprint available at arXiv:2510.19718.
- [24] R. Hickingbotham, D. Y. Kang, S.-I. Oum, R. Steiner, and D. R. Wood, Clustered colouring of odd- H -minor-free graphs, in *2023 MATRIX annals*, *MATRIX Book Ser.*, vol. 6, Springer, Cham, 2025, 73–80.
- [25] S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000.
- [26] T. R. Jensen and B. Toft, *Graph coloring problems*, Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., New York, 1995.
- [27] Y. Ji, Z.-X. Song, E. Weiss, and X. Zhang, Odd clique minors in graphs with independence number two, 2025. Preprint available at arXiv:2505.07727.
- [28] D. Y. Kang and S.-I. Oum, Improper colouring of graphs with no odd clique minor, *Combin. Probab. Comput.* **28** (2019), 740–754.
- [29] K.-i. Kawarabayashi, A weakening of the odd Hadwiger's conjecture, *Combin. Probab. Comput.* **17** (2008), 815–821.

- [30] K.-i. Kawarabayashi, Note on coloring graphs without odd- K_k -minors, *J. Combin. Theory Ser. B* **99** (2009), 728–731.
- [31] K.-i. Kawarabayashi and B. Reed, Fractional coloring and the odd Hadwiger’s conjecture, *European J. Combin.* **29** (2008), 411–417.
- [32] K.-i. Kawarabayashi and Z.-X. Song, Some remarks on the odd Hadwiger’s conjecture, *Combinatorica* **27** (2007), 429–438.
- [33] J. H. Kim, The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$, *Random Structures Algorithms* **7** (1995), 173–207.
- [34] A. V. Kostochka, The minimum Hadwiger number for graphs with a given mean degree of vertices, *Metody Diskret. Analiz.* (1982), 37–58.
- [35] A. V. Kostochka, Lower bound of the Hadwiger number of graphs by their average degree, *Combinatorica* **4** (1984), 307–316.
- [36] M. Krivelevich, Bounding Ramsey numbers through large deviation inequalities, *Random Structures Algorithms* **7** (1995), 145–155.
- [37] C. McDiarmid and B. Reed, Concentration for self-bounding functions and an inequality of Talagrand, *Random Structures Algorithms* **29** (2006), 549–557.
- [38] J. Nešetřil and R. Thomas, A note on spatial representation of graphs, *Comment. Math. Univ. Carolin.* **26** (1985), 655–659.
- [39] S. Norin, Recent progress towards Hadwiger’s conjecture, in *ICM—International Congress of Mathematicians. Vol. 6. Sections 12–14*, EMS Press, Berlin, 2023, 4606–4620.
- [40] S. Norin, L. Postle, and Z.-X. Song, Breaking the degeneracy barrier for coloring graphs with no K_t minor, *Adv. Math.* **422** (2023), Paper No. 109020, 23pp.
- [41] S. Norin and Z.-X. Song, A new upper bound on the chromatic number of graphs with no odd K_t minor, *Combinatorica* **42** (2022), 137–149.
- [42] L. Postle, Further progress towards the list and odd versions of Hadwiger’s conjecture, 2020. Preprint available at arXiv:2010.05999.
- [43] G. Ringel and J. W. T. Youngs, Solution of the Heawood map-coloring problem, *Proc. Nat. Acad. Sci. U.S.A.* **60** (1968), 438–445.
- [44] N. Robertson, P. Seymour, and R. Thomas, Hadwiger’s conjecture for K_6 -free graphs, *Combinatorica* **13** (1993), 279–361.
- [45] P. Seymour, Hadwiger’s conjecture, in *Open problems in mathematics*, Springer, Cham, 2016, 417–437.
- [46] J. B. Shearer, A note on the independence number of triangle-free graphs, *Discrete Math.* **46** (1983), 83–87.
- [47] G. Simonyi and A. Zsbán, On topological relaxations of chromatic conjectures, *European J. Combin.* **31** (2010), 2110–2119.
- [48] J. Spencer, Asymptotic lower bounds for Ramsey functions, *Discrete Math.* **20** (1977/78), 69–76.
- [49] R. Steiner, Asymptotic equivalence of Hadwiger’s conjecture and its odd minor-variant, *J. Combin. Theory Ser. B* **155** (2022), 45–51.
- [50] R. Steiner, Improved bound for improper colourings of graphs with no odd clique minor, *Combin. Probab. Comput.* **32** (2023), 326–333.
- [51] R. Steiner, Hadwiger’s conjecture and topological bounds, *European J. Combin.* **122** (2024), Paper No. 104033, 11pp.
- [52] M. Talagrand, Concentration of measure and isoperimetric inequalities in product spaces, *Inst. Hautes Études Sci. Publ. Math.* (1995), 73–205.
- [53] A. Thomason, An extremal function for contractions of graphs, *Math. Proc. Cambridge Philos. Soc.* **95** (1984), 261–265.

- [54] C. Thomassen, Graph decomposition with applications to subdivisions and path systems modulo k , *J. Graph Theory* **7** (1983), 261–271.
- [55] C. Thomassen, Totally odd K_4 -subdivisions in 4-chromatic graphs, *Combinatorica* **21** (2001), 417–443.
- [56] B. Toft, A survey of Hadwiger’s conjecture, *Congr. Numer.* **115** (1996), 249–283.
- [57] Y. Wang, Improved bound for Hadwiger’s conjecture, *Graphs Combin.* **41** (2025), Paper No. 36, 13pp.
- [58] W. Zang, Proof of Toft’s conjecture: every graph containing no fully odd K_4 is 3-colorable, *J. Comb. Optim.* **2** (1998), 117–188.

APPENDIX A. OPTIMALITY OF THE FACTOR $\frac{3}{2}$

In this section, we prove Proposition 1.5, whose proof is inspired by similar arguments used in [27]. In the proof, we make use of the following result of Gallai [19]. Recall that a graph is called k -critical if it has chromatic number k but every proper subgraph is $(k - 1)$ -colorable.

Theorem A.1. *Let $k \geq 2$ be an integer, and let G be a k -critical graph. If $|V(G)| \leq 2k - 2$, then there exists a partition $V(G) = X_1 \cup X_2$ of $V(G)$ into two non-empty sets X_1 and X_2 such $\{x_1, x_2\} \in E(G)$ for all $x_1 \in X_1$ and $x_2 \in X_2$.*

With this tool, we are ready to present the proof of Proposition 1.5.

Proof of Proposition 1.5. We prove the statement in the proposition by contradiction. Supposing the statement is false, let t and G be chosen to form a counterexample minimizing $n := |V(G)|$. Then $\alpha(G) \leq 2$, the graph G does not contain K_t as an odd minor, and we have $k := \chi(G) > \lceil \frac{3}{2}(t - 1) \rceil$ and therefore $k - 1 \geq \lceil \frac{3}{2}(t - 1) \rceil$. By [32, Theorem 1.8], G contains $K_{\lceil n/3 \rceil}$ as an odd minor, so we have $\lceil n/3 \rceil \leq t - 1$ and hence $n \leq 3t - 3 \leq 2k - 2$. Observe that by our minimality assumption on $|V(G)|$, we have $\chi(G - v) = k - 1$ for every $v \in V(G)$. Letting G' be a k -critical subgraph of G , we must therefore have $V(G') = V(G)$, and by Theorem A.1 we find that there exists a partition $V(G) = V(G') = X_1 \cup X_2$ of $V(G') = V(G)$ into non-empty subsets X_1 and X_2 such that all possible edges between X_1 and X_2 exist in G' and thus also in G .

Let $G_1 := G[X_1]$ and $G_2 := G[X_2]$, and note that we have $\chi(G) = \chi(G_1) + \chi(G_2)$. For $i \in \{1, 2\}$, let t_i denote the largest positive integer such that G_i has K_{t_i} as an odd minor. By our minimality assumption on $|V(G)|$, we then have $\chi(G_i) \leq \lceil \frac{3}{2}t_i \rceil$ for $i \in \{1, 2\}$.

Finally, for $i \in \{1, 2\}$, let $Y_i \subseteq X_i$ be chosen as an inclusion-wise minimal subset of X_i such that $G[Y_i]$ still contains K_{t_i} as an odd minor, and let $T_1^i, \dots, T_{t_i}^i$ be a collection of vertex-disjoint trees in $G[Y_i]$ as in the definition of odd minors. This means that for $i \in \{1, 2\}$ there exist proper colorings of the vertices of each of the trees $T_1^i, \dots, T_{t_i}^i$ with the colors black and white, such that any two of these trees are connected by some monochromatic edge. Note that each of these trees is either a single vertex or has both black and white vertices, and that for each $i \in \{1, 2\}$ all single-vertex trees among $T_1^i, \dots, T_{t_i}^i$ must have the same color. Swapping the colors black and white if needed in the colorings of $T_1^i, \dots, T_{t_i}^i$ for one or both $i \in \{1, 2\}$, we may assume without loss of generality that for each $i \in \{1, 2\}$, all trees $T_1^i, \dots, T_{t_i}^i$ contain a white vertex.

Now, by considering the collection of trees $T_1^1, \dots, T_{t_1}^1, T_1^2, \dots, T_{t_2}^2$ we find $K_{t_1+t_2}$ as an odd minor in G . This means that $t_1 + t_2 \leq t - 1$. Moving on, we distinguish two cases.

Case 1: There exist $i \in \{1, 2\}$ and $y \in Y_i$ such that $(X_i \setminus Y_i) \cup \{y\}$ is independent.

Let us assume without loss of generality that $i = 1$ (the case $i = 2$ is analogous). In this case, we can obtain a better estimate on the chromatic number of G_1 , as follows. By minimality of Y_1 , the graph $G'_1 := G_1 - ((X_1 \setminus Y_1) \cup \{y\}) = G[Y_1 \setminus \{y\}]$ does not contain K_{t_1} as an odd minor. Thus,

since $\alpha(G'_1) \leq \alpha(G) \leq 2$, and by minimality of G , we obtain $\chi(G_1) \leq \chi(G'_1) + 1 \leq \lceil \frac{3}{2}(t_1 - 1) \rceil + 1$. We therefore have

$$\chi(G) = \chi(G_1) + \chi(G_2) \leq \left(\left\lceil \frac{3(t_1 - 1)}{2} \right\rceil + 1 \right) + \left\lceil \frac{3t_2}{2} \right\rceil = \left\lceil \frac{3(t_1 - 1)}{2} \right\rceil + \left\lceil \frac{3t_2}{2} \right\rceil + 1.$$

Consider first the case that t_1 is even and t_2 is odd. Then we have $\lceil \frac{3(t_1 - 1)}{2} \rceil = \frac{3(t_1 - 1)}{2} + \frac{1}{2}$ and $\lceil \frac{3t_2}{2} \rceil = \frac{3t_2}{2} + \frac{1}{2}$, as well as $\lceil \frac{3(t_1 + t_2)}{2} \rceil = \frac{3(t_1 + t_2)}{2} + \frac{1}{2}$ (since $t_1 + t_2$ is odd). So we can conclude

$$\chi(G) \leq \left(\frac{3(t_1 - 1)}{2} + \frac{1}{2} \right) + \left(\frac{3t_2}{2} + \frac{1}{2} \right) + 1 = \frac{3(t_1 + t_2)}{2} + \frac{1}{2} = \left\lceil \frac{3(t_1 + t_2)}{2} \right\rceil \leq \left\lceil \frac{3(t - 1)}{2} \right\rceil.$$

In all other cases, we have that t_1 is odd (and hence $t_1 - 1$ is even) or t_2 is even and thus $\lceil \frac{3(t_1 - 1)}{2} \rceil + \lceil \frac{3t_2}{2} \rceil \leq \frac{3(t_1 - 1)}{2} + \frac{3t_2}{2} + \frac{1}{2}$ holds, since the expressions $\lceil \frac{3(t_1 - 1)}{2} \rceil - \frac{3(t_1 - 1)}{2}$ and $\lceil \frac{3t_2}{2} \rceil - \frac{3t_2}{2}$ are both bounded by $\frac{1}{2}$ and at least one of these two expressions is equal to 0. Therefore we obtain

$$\chi(G) \leq \left\lceil \frac{3(t_1 - 1)}{2} \right\rceil + \left\lceil \frac{3t_2}{2} \right\rceil + 1 \leq \frac{3(t_1 - 1)}{2} + \frac{3t_2}{2} + \frac{1}{2} + 1 = \frac{3(t_1 + t_2)}{2} \leq \left\lceil \frac{3(t - 1)}{2} \right\rceil.$$

In any case, this yields a contradiction to the assumption $\chi(G) > \lceil \frac{3}{2}(t - 1) \rceil$ and concludes the proof of Case 1.

Case 2: For all $i \in \{1, 2\}$ and $y \in Y_i$, the set $(X_i \setminus Y_i) \cup \{y\}$ is not independent.

This, in particular, implies that $X_1 \setminus Y_1 \neq \emptyset$ and $X_2 \setminus Y_2 \neq \emptyset$. Let $T \subseteq G[(X_1 \setminus Y_1) \cup (X_2 \setminus Y_2)]$ be a tree with a proper 2-coloring defined as follows. If $X_1 \setminus Y_1$ is an independent set in G , then we let T be any spanning tree of the complete bipartite graph spanned between $X_1 \setminus Y_1$ and $X_2 \setminus Y_2$ in G , and we color all vertices in $X_1 \setminus Y_1$ white and all vertices in $X_2 \setminus Y_2$ black. On the other hand, if the set $X_1 \setminus Y_1$ is not independent, then we pick two adjacent vertices $u, v \in X_1 \setminus Y_1$, and some vertex $w \in X_2 \setminus Y_2$, and define T as the path with vertices u, v, w in G (in this order), with u and w colored white and v colored black.

We claim that now, every vertex in $Y_1 \cup Y_2$ has a white vertex in T as a neighbor. To verify this, consider first the case that $X_1 \setminus Y_1$ is independent, then the set of white vertices in T is $X_1 \setminus Y_1$. Now, by assumption of Case 2, every vertex $y \in Y_1$ must have a neighbor in $X_1 \setminus Y_1$ (i.e., a white neighbor in T). Furthermore, every vertex $y \in Y_2 \subseteq X_2$ also has a neighbor in $X_1 \setminus Y_1$ (since all possible edges between X_1 and X_2 exist in G). Thus, every vertex in $Y_1 \cup Y_2$ has a white vertex in T as a neighbor if $X_1 \setminus Y_1$ is independent.

Next, consider the second case, where T is the path on vertices u, v, w with $u, v \in X_1$ and $w \in X_2$, and u, w are the white vertices of T . Now, every vertex $y \in Y_1 \subseteq X_1$ is a neighbor of w and every vertex $y \in Y_2 \subseteq X_2$ is a neighbor of u . Thus, again every vertex in $Y_1 \cup Y_2$ has a white vertex in T as a neighbor.

Now, the collection of the trees $T_1^1, \dots, T_{t_1}^1, T_1^2, \dots, T_{t_2}^2, T$ shows that G contains $K_{t_1 + t_2 + 1}$ as an odd minor in G . Indeed, we clearly have a monochromatic edge between any two of the trees $T_1^1, \dots, T_{t_1}^1$ as well as between any two of the trees $T_1^2, \dots, T_{t_2}^2$, by how we chose these trees. Furthermore, for any $j_1 \in [t_1]$ and $j_2 \in [t_2]$ there is a monochromatic edge between the trees $T_{j_1}^1$ and $T_{j_2}^2$, since both of them contain a white vertex, and these two white vertices form an edge in G (since G contains all possible edges between X_1 and X_2). Finally, every vertex in $Y_1 \cup Y_2$ has a white neighbor in T , so in particular from each of the trees $T_1^1, \dots, T_{t_1}^1, T_1^2, \dots, T_{t_2}^2$ there is a monochromatic edge to T .

Since we assumed that G does not contain K_t as an odd minor, we can now conclude that $t_1 + t_2 + 1 \leq t - 1$. This implies

$$\chi(G) = \chi(G_1) + \chi(G_2) \leq \left\lceil \frac{3t_1}{2} \right\rceil + \left\lceil \frac{3t_2}{2} \right\rceil \leq \left(\frac{3t_1}{2} + \frac{1}{2} \right) + \left(\frac{3t_2}{2} + \frac{1}{2} \right) < \frac{3(t_1 + t_2 + 1)}{2} \leq \frac{3}{2}(t-1),$$

contradicting our initial assumption $\chi(G) > \lceil \frac{3}{2}(t-1) \rceil$. This concludes the proof of Case 2 and thereby of the whole proposition. \square