Cubes, Grids and Induced Subgraphs: Topics in Extremal and Probabilistic Combinatorics



Michael Savery

Lincoln College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2024

If thou dislik'st the piece thou light'st on first;
Think that of all that I have writ, the worst.

From To the Sour Reader,
Robert Herrick

Abstract

We begin by studying random subgraphs with a certain type of local dependence: a bond percolation model on a graph G is said to be 1-independent if for all vertex-disjoint sets of edges S_1 and S_2, the states of the edges in S_1 are independent of the states of the edges in S_2. In 2012, Balister and Bollobás defined $p_{\max}(G)$ to be the supremum of those p for which there exists a 1-independent bond percolation model on G in which each edge is present in the random subgraph with probability at least p, but which does not percolate (i.e. which does not with positive probability produce an infinite connected component). A fundamental and challenging problem in this area is to determine the value of $p_{\max}(G)$ when G is the lattice graph $\mathbb{Z}^2$. Since $p_{\max}(\mathbb{Z}^n) \leqslant p_{\max}(\mathbb{Z}^{n-1})$, it is also of interest to establish the value of $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$. We improve the known upper bounds on both quantities, the latter significantly. In proving these results, we also give an upper bound on the critical probability for a 1-independent model on the hypercube graph to contain a giant component.

Next, we analyse the typical structure of multi-agent strategic games in terms of the connectivity properties of their best-response graphs. Our central result shows that almost every game that is generic (that is, without indifferences) and has a pure Nash equilibrium and a 'large' number of players is 'connected', meaning that every action profile that is not a pure Nash equilibrium can reach every pure Nash equilibrium via a best-response path. This has important implications for dynamics in games. In particular, we show that there are simple, uncoupled, adaptive dynamics for which period-by-period play converges almost surely to a pure Nash equilibrium in almost every large generic game that has one. This contrasts with the known fact that there is no such dynamic that leads almost surely to a pure Nash equilibrium in *every* generic game that has one. For our characterisation of game connectivity we build on recent work of McDiarmid, Scott, and Withers on the component structure of dense random subgraphs of the hypercube.

Following this, we prove some results on the relationship between a graph's chromatic number and its induced subgraphs. We show that, for every graph F with at least one edge, there is a constant c_F such that there are graphs of arbitrarily large chromatic number and the same clique number as F in which every F-free induced subgraph has chromatic number at most c_F. This has important consequences for the study of χ-bounded graph classes and generalises recent theorems of Briański, Davies, and Walczak, and Carbonero, Hompe, Moore, and Spirkl. Our results imply that for every $r \geqslant 3$ the class of K_r-free graphs has a very strong vertex Ramsey-type property, giving a generalisation of a result of Folkman from 1970. We also prove related results for tournaments, hypergraphs, and infinite families of graphs, and show an analogous statement for graphs where clique number is replaced by odd girth.

Finally, we consider the problem of maximising the number of subgraphs of an n-vertex planar graph that are isomorphic to some fixed graph H. We present results in two settings: that in which the copies of H are required to be induced and that in which they are not. In the induced setting, we determine the n-vertex extremal graphs for large n when H is the paw graph, and for all n when H is K_4^- or a star. In the not necessarily induced setting we determine the n-vertex extremal graphs for all n when H is the paw graph or K_4^-, with the cases where H is a star having been solved by Alon and Caro.

Acknowledgements

Thank you to Alex for guiding me through the world of research mathematics over the past four years. Thank you for proposing interesting problems, for sage advice on all things, for always having time for a chat, and for fostering such a friendly, collaborative, and supportive group.

Thank you to my examiners, Imre Leader and Oliver Riordan, for your comments, corrections, and improvements, and for catching a few howlers.

Thank you to Julia Gog, Thomas Forster, and Thomas Bloom for your guidance and encouragement during my undergraduate studies, and thank you to the Heilbronn Institute for Mathematical Research for supporting me through this DPhil. Collaboration has been a key theme of my time in Oxford and has been something I've greatly enjoyed: thank you to all my collaborators for the time and effort you've invested.

Going into the department has rarely felt like going into work,[1] and this owes much to the people to be found there. It's been such a joy to have shared rotis, whiteboards, croquet lawns, sing-alongs, Overcooked kitchens, walks, trips, puzzles, games nights, two-hour lunches, and baffled mid-seminar glances with Emil, Jane, Tom, Youri, Freddie, António, Lukas, Emma, and Kate. Thank you for making me so sad to leave.

Thank you also to all the other Oxford combinatorialists past and present who have made the group such a pleasure to be part of.

Thank you to my old friends for continuing to put up with me. Ben and Ben of course deserve distinguished thanks on this front, and special thanks go also to Alex, Mark, and Thomas for regularly checking in.

Finally, thank you to my family for your continued love and support, and for always being at the end of the phone.

[1] Wisecrack left as an exercise to the reader.

Statement of Originality

The material in this thesis has not been submitted for any other degree or qualification at any university.

Chapter 2 is based on joint work with Paul Balister, Tom Johnston, and Alex Scott [18]. Chapter 3 is based on joint work with Tom Johnston, Alex Scott, and Bassel Tarbush [87]. Chapter 4 is based on joint work with António Girão, Freddie Illingworth, Emil Powierski, Alex Scott, Youri Tamitegama, and Jane Tan [62].

Contents

1 Introduction **1**

1.1 Locally dependent percolation 1

1.2 Game connectivity and adaptive dynamics 5

1.3 Induced subgraphs of induced subgraphs of large chromatic number . 10

1.4 Planar inducibility 13

2 Locally dependent percolation on $\mathbb{Z}^n$ **18**

2.1 Notation and preliminaries 20

2.2 Giant component in Q_n 21

2.2.1 Proof of Theorem 2.3 22

2.2.2 Proof of Proposition 2.7 25

2.3 Percolation in $\mathbb{Z}^n$ 33

2.4 Extracting a concrete bound 35

2.5 Upper bounds on $p_{\max}(\mathbb{Z}^2)$ 37

2.6 Open problems 45

3 Game connectivity and adaptive dynamics **50**

3.1 Game connectivity 54

3.2 Directed grids 55

3.2.1 Game connectivity 55

3.2.2 Relation to previous work 58

3.2.3 Outline of the proof of Theorem 3.8 59

3.2.4 Acyclicity of directed grids 60

3.3 Proof of Theorem 3.8 61

3.3.1 Establishing a foothold 62

3.3.2 A strongly connected component 66

3.3.3 Accessing good vertices 71

3.4 The tightness of Theorems 3.5 and 3.8 73

3.5 Proof of Proposition 3.12 . . . 78
3.6 Adaptive dynamics and equilibrium selection . . . 80
3.7 Open questions . . . 85

4 Induced subgraphs of induced subgraphs of large chromatic number 88
4.1 Orientations of classical constructions . . . 91
4.2 Clique number . . . 93
4.2.1 B_3-sets . . . 94
4.2.2 Proof of Theorem 1.8 . . . 96
4.3 Odd girth . . . 98
4.3.1 B_h-sets . . . 98
4.3.2 Proof of Theorem 4.1 . . . 100
4.4 Tournaments . . . 103
4.5 Hypergraphs . . . 104
4.6 Extensions to infinite families of graphs . . . 109
4.7 Conclusion . . . 114

5 Planar inducibility 116
5.1 Preliminaries . . . 121
5.1.1 Notation and definitions . . . 121
5.1.2 Preliminary results . . . 122
5.1.3 A note on computer checks . . . 124
5.2 Induced paws . . . 124
5.3 Induced copies of K_4^- . . . 136
5.4 Induced stars . . . 139
5.4.1 Small stars . . . 139
5.4.2 Big stars . . . 142
5.5 Results on $f(n, H)$ . . . 145
5.5.1 Paws . . . 145
5.5.2 K_4^- . . . 147
5.6 Some corollaries of the main theorems . . . 149
5.7 Concluding remarks . . . 152

Bibliography 154

List of Figures

1.1 Best-response graph of a generic, connected, but not acyclic three-player, two-action game. The • vertices are sinks. The ∘ vertices form a cycle. 6

2.1 The sets $S_u^{r,s}$ 38
2.2 A portion of the graph $\mathbb{Z}^2$ with the regions corresponding to the $\mathcal{L}_u$ when $\theta = 0$ highlighted. Dashed edges represent edges open with probability p', other edges are open with probability p. 39
2.3 The configurations corresponding to the definitions of $\mathcal{C}_0$, $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$ 39
2.4 A portion of the graph $\mathbb{Z}^2$ and the line $x + y = 1$. Dashed edges represent edges open with probability at least p', other edges are open with probability at least p. 43
2.5 A portion of the reflected $(k-1)$-times renormalised model. If F_{k-1} holds, then the edge $(1,0)(1,1)$ is open, so if there is a component intersecting both $S_{(0,0)}$ and $S_{(1,0)}$ in large sets, then this component contains that edge. 44

5.1 The exceptional graphs for Theorems 5.2 and 5.4 118
5.2 The n-vertex planar graphs containing the most induced copies of K_4^- for small n 119
5.3 An illustration of part of the graph G' in Case 2(i) 126
5.4 The extra paw induced on $\{v, u_1, x, z\}$ that gives a contradiction in Case 2(ii). The dashed lines represent other edges incident to u_1 that may or may not be present 127
5.5 An illustration of part of the drawing of G_1 130
5.6 The only non-extremal n-vertex graph with more than $g(n)$ induced copies of K_4^- 137
5.7 A planar graph containing many induced copies of P_4 152

Chapter 1

Introduction

This thesis addresses various problems arising in extremal and structural graph theory and in probabilistic combinatorics. These can broadly be partitioned into two groups: Chapters 2 and 3 concern problems on random subgraphs of certain cube- or grid-like graphs, while the work in Chapters 4 and 5 centres around induced subgraphs. In this chapter we summarise, motivate, and put in context the results that follow.

1.1 Locally dependent percolation

A *percolation model* on a (possibly infinite) graph G is a probability distribution on the subgraphs of G. We say that a vertex or edge of G is *open* if it is present in the random subgraph of G associated with the percolation model, and *closed* otherwise. A *bond percolation model* on G is a probability distribution on the spanning subgraphs of G, meaning all vertices are open and the random subgraph depends only on the edges. An *independent* bond percolation model on G is a bond percolation model in which each edge is open independently.

The focus of Chapter 2 is on the weaker condition of 1-independence: a bond percolation model on G is said to be *1-independent* if, for any two disjoint sets of edges S_1 and S_2 of G such that no edge in S_1 shares a vertex with an edge in S_2, the states (i.e. open or closed) of the edges in S_1 are independent of the states of the edges in S_2. If G is an infinite graph, then we say that a percolation model on G *percolates* if the associated random subgraph of G has an infinite component with positive probability. Note that by a slightly modified version of Kolmogorov's zero-one law, if a percolation model on a locally-finite infinite graph G percolates, then in fact the random subgraph has an infinite component with probability 1.

The study of 1-independent bond percolation models is motivated by their use as a tool to obtain bounds on critical probabilities of independent models via renormalisation (see, for example, [17] and Sections 3.5 and 6.2 of [24]). They have also been combined with renormalisation techniques to analyse bond percolation models with dependencies over a greater range (see [17] again for an example).

For an infinite graph G, we wish to obtain results on whether or not a 1-independent model on G in which each edge is open with some probability $p \in (0,1)$ percolates. For any given p there will be many different 1-independent models, but we might expect that if p is sufficiently large, then any 1-independent model on G will percolate. In this spirit, Balister and Bollobás [16] defined $\mathcal{D}_{\geqslant p}(G)$ to be the class of 1-independent bond percolation models on G in which each edge is open with probability at least p. They then defined

$$p_{\max}(G) \coloneqq \sup\big\{p : \text{some model in } \mathcal{D}_{\geqslant p}(G) \text{ does not percolate}\big\}.$$

Note that in this definition it is equivalent to ask for a model in which each edge is open with probability exactly p. Indeed, if $\mu \in \mathcal{D}_{\geqslant p}(G)$ does not percolate, then the model in which we sample from μ and then independently delete each edge e with probability $1 - p/\mathbb{P}_\mu(e \text{ open})$ is a 1-independent bond percolation model which does not percolate, and in which each edge is open with probability exactly p.

The value of $p_{\max}(G)$ for different infinite graphs G has been studied, along with various other properties of 1-independent bond percolation models, in [16, 17, 40, 51, 99]. In particular, in [16] the authors determine $p_{\max}(T)$ for all locally finite infinite trees T in terms of a parameter known as the *branching number* of T. They go on to establish that $p_{\max}(G) \geqslant \frac{1}{2}$ for all locally finite connected infinite graphs G, and they construct such a G which achieves this lower bound.

The case where G is the lattice graph $\mathbb{Z}^2$ has been of particular interest, with Balister and Bollobás asking for the value of $p_{\max}(\mathbb{Z}^2)$.

Question 1.1 ([16])**.** *What is the value of* $p_{\max}(\mathbb{Z}^2)$*?*

Considering $\mathbb{Z}^n$ more generally, they noted that $p_{\max}(\mathbb{Z}^n)$ is decreasing in n (since if a model does not percolate on $\mathbb{Z}^n$ then it does not percolate on any $\mathbb{Z}^{n-1}$ subspace), and asked for the limit of this value as $n \to \infty$.

Question 1.2 ([16])**.** *What is the value of* $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$*?*

The best known lower bound on $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$ comes from a construction of Day, Falgas-Ravry, and Hancock [40] which gives

$$p_{\max}(\mathbb{Z}^n) \geqslant 4 - 2\sqrt{3} = 0.535898\ldots$$

for all $n \geqslant 2$. It was conjectured by Falgas-Ravry and Pfenninger [51] that there exists some $n \geqslant 3$ for which this is an equality. In [40] the authors also give a separate construction which yields

$$p_{\max}(\mathbb{Z}^2) \geqslant p_{\text{site}}^2 + \tfrac{1}{2}(1 - p_{\text{site}}),$$

where $p_{\text{site}} = p_{\text{site}}(\mathbb{Z}^2)$ denotes the critical probability for independent site percolation on $\mathbb{Z}^2$, that is, the supremum of the p such that percolation does not occur in the percolation model on $\mathbb{Z}^2$ in which vertices are open independently with probability p and edges are open if both their endpoints are open. The best known rigorous lower bound on p_{site} is 0.556 due to van den Berg and Ermakov [143]. Substituting this into the above gives $p_{\max}(\mathbb{Z}^2) \geqslant 0.531136$ which is slightly worse than the bound given by the first construction. However, using the non-rigorous estimate $p_{\text{site}} \approx 0.592746$ (see, for example, [157]) gives an improved bound of $p_{\max}(\mathbb{Z}^2) \geqslant 0.554974$.

In [18], with Paul Balister, Tom Johnston, and Alex Scott we improved the rigorous lower bound on $p_{\max}(\mathbb{Z}^2)$ to $p_{\max}(\mathbb{Z}^2) \geqslant \frac{1}{32}(35 - 3\sqrt{33}) = 0.555197\ldots$, which beats even the non-rigorous bound stated above. However, it is likely that the true value of $p_{\max}(\mathbb{Z}^2)$ is still some distance from this, as a high confidence result (p-value $< 10^{-11}$) in [18] suggests that in fact $p_{\max}(\mathbb{Z}^2) \geqslant 0.5921$.

The previous best known upper bound on $p_{\max}(\mathbb{Z}^2)$ was 0.8639 due to Balister, Bollobás, and Walters [17], and no better upper bounds on $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$ were known. In Chapter 2, we improve both upper bounds as follows.

Theorem 1.3. $p_{\max}(\mathbb{Z}^2) \leqslant 0.8457$.

Theorem 1.4. $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n) \leqslant 0.5847$.

The proof of Theorem 1.3 uses a 2×2 renormalisation argument similar to that employed in [17]. We obtain a better bound than was derived in that paper by using a different condition for two renormalised sites to be joined by an edge. To calculate a lower bound on the probability that each edge is present, we solve two linear programming problems that are relaxations of the 1-independence constraints on a 4×2 grid. In Section 2.5 we build on this renormalisation to show that under every model in $\mathcal{D}_{\geqslant 0.8457}(\mathbb{Z}^2)$ there are almost surely open cycles enclosing any bounded

region of the plane (so that, in particular, almost surely the infinite component is unique and there is no infinite component in the dual graph), and that under every model in $\mathcal{D}_{\geqslant 0.8459}(\mathbb{Z}^2)$ the origin is in an infinite component with positive probability.

We denote by Q_n the n-dimensional *hypercube graph*, which has vertex set $\{0,1\}^n$ and in which vertices are joined by an edge if they differ in exactly one coordinate. The proof of Theorem 1.4 relies on a result (Theorem 2.1) which translates a condition on 1-independent bond percolation models on hypercube graphs to one on the lattice. To apply this result we solve two more linear programming problems, this time relaxations of the 1-independence constraints on Q_3.

The approach we use to prove Theorem 2.1 also yields a result about the critical probability for the asymptotically almost sure existence of a giant component in the hypercube graph Q_n under 1-independent bond percolation models, which may be of independent interest. To this end, we let

$$p_{\text{giant}} \coloneqq \inf\left\{p : \exists C > 0 : \lim_{n\to\infty} \inf_{\mu\in\mathcal{D}_{\geqslant p}(Q_n)} \mathbb{P}_\mu(Q_n \text{ has component of size} \geqslant C\cdot 2^n) = 1\right\}$$

be the threshold p for the asymptotically almost sure existence of a giant component in Q_n under any model in $\mathcal{D}_{\geqslant p}(Q_n)$. We obtain the same upper bound on p_{giant} as we do on $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$.

Theorem 1.5. $p_{\text{giant}} \leqslant 0.5847$.

The best known lower bound on p_{giant} comes from a simple construction used in [16], which shows that $p_{\text{giant}} \geqslant \frac{1}{2}$. It was conjectured by Falgas-Ravry and Pfenninger [51] that this lower bound is tight, that is, $p_{\text{giant}} = \frac{1}{2}$.

It is interesting to compare the known bounds on these 1-independent threshold probabilities with their analogues in the independent setting. First, the celebrated Harris–Kesten theorem [71, 91] determines the threshold for percolation in the independent bond model on $\mathbb{Z}^2$ to be exactly $1/2$, whereas the high-confidence result in [18] suggests that $p_{\max}(\mathbb{Z}^2) \geqslant 0.5921$. Turning to p_{giant}, Erdős and Spencer showed in [46] that for each fixed $\varepsilon > 0$, the random subgraph of Q_n in which each edge is open independently with probability $p = (1-\varepsilon)/n$ has maximum component size $o(2^n)$ asymptotically almost surely, while Ajtai, Komlós, and Szemerédi showed in [1] that under the same model with $p = (1+\varepsilon)/n$ where $\varepsilon > 0$ is fixed, there is a component of size $\Omega(2^n)$. The component structure has since been studied when $\varepsilon \to 0$ at different rates, for example in [23, 25–27, 145, 146] (see also [144] or Section 13.4 of [82] for a survey). The threshold probability for percolation of the independent

model on $\mathbb{Z}^n$ is $(1 + o(1))/2n$; see [82] (particularly Section 15.5) for the history of this fact and a survey of more precise results.

Chapter 2 is based on joint work with Paul Balister, Tom Johnston, and Alex Scott [18].

1.2 Game connectivity and adaptive dynamics

A fundamental question at the heart of the literature on learning in games and distributed systems is whether there are simple dynamics that are guaranteed to lead to a Nash equilibrium in every game (see e.g. [57, 153]). Several influential results have outlined the boundary between the possible and the impossible, i.e. between classes of dynamics that are guaranteed to lead to a Nash equilibrium in every game and classes that lack such a guarantee (see [154] for a survey). It has been shown, in particular, that there are no 'simple' adaptive dynamics that lead to a pure Nash equilibrium in every game that has one [13, 76–78, 86]. In Chapter 3, we investigate how this boundary shifts when the requirement of convergence in *every* game is relaxed to *almost every* game.

We approach this problem as follows. The behaviours of many game dynamics are crucially determined by the connectivity properties of a game's best-response graph.[1] We therefore classify games according to such properties, and build on recent combinatorial results of McDiarmid, Scott, and Withers [105] to quantify the relative sizes of these classes. This approach differs significantly from much of the literature on adaptive dynamics, where the properties of the dynamics themselves are the more usual focus.

The central game connectivity result of Chapter 3 is that the best-response graph of almost every game that is 'generic' (i.e. has no indifferences), that has a pure Nash equilibrium, and that has a large number of players relative to the number of actions available to each player, has the property that every action profile that is not a pure Nash equilibrium can reach *every* pure Nash equilibrium profile via best-response paths. In sum, almost every large generic game that has a pure Nash equilibrium satisfies a strong connectivity property.

Moreover, this property is conducive to convergent dynamics: combining our result on game connectivity with existing results on the convergence properties of various

[1] See Chapter 3 for definitions not given here.

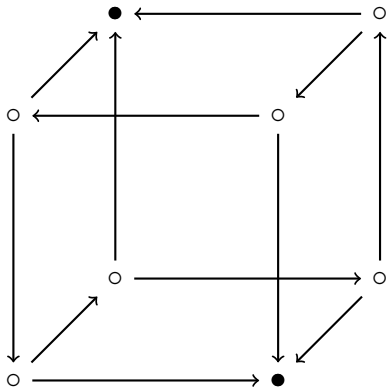


Figure 1.1: Best-response graph of a generic, connected, but not acyclic three-player, two-action game. The • vertices are sinks. The ∘ vertices form a cycle.

game dynamics implies that there exist simple adaptive dynamics that lead to a pure Nash equilibrium in *almost every* large generic game that has one, even though there are no such dynamics for which this holds in *every* such game.

Game connectivity The *best-response graph* of a game is the directed graph whose vertex set is the set of pure action profiles of the game, and whose directed edges correspond to best-responses [152]. A game's pure Nash equilibria correspond to the sinks of its best-response graph. Games can be classified according to the connectivity properties of their best-response graphs: well known classes include *weakly acyclic* games (games whose best-response graphs have the property that every vertex can reach a sink along a directed path) and *acyclic* games (games whose best-response graphs contain no cycles). Clearly every acyclic game is weakly acyclic but the converse does not hold (indeed, acyclicity is a much more restrictive than weak acyclicity).

We introduce a new class of games which we refer to as connected games. A game is *connected* if its best-response graph has a sink and has the property that every non-sink can reach *every* sink along a directed path. Connectedness is a strictly stronger property than weak acyclicity. The three-player game with two actions per player whose best-response graph is illustrated in Figure 1.1 is connected: there are two sinks and every non-sink can reach every sink along a directed best-response path. However, if the downwards edge on the left were flipped to point upwards, then the game would no longer be connected, although it would still be weakly acyclic.

We take players' preferences to be encoded ordinally via a preference relation rather than cardinally via a utility function, and, throughout, we focus on games that are

generic, by which we mean that there are no indifferences. The game whose best-response graph is illustrated in Figure 1.1, for example, is generic.

The central game connectivity result of Chapter 3 (Theorem 3.5) is that:

Almost every large generic game that has a pure Nash equilibrium is connected.

Here, by 'large' we mean that the number of players is much greater than the maximum number of actions available to each player, and by 'almost every' we mean all but an exponentially small proportion (in terms of the number of players). This result allows for different numbers of actions to be available to each player, provided the maximum number of actions is small relative to the number of players.

The result above follows from stronger results regarding the connectivity properties of random subgraphs of directed Hamming graphs, which we present in Section 3.2. These are the main technical contribution of our work. The proof of our main technical result (Theorem 3.8) adapts new work of McDiarmid, Scott, and Withers [105] on the component structure of random subgraphs of the hypercube.

The recent work of Amiet, Collevecchio, Scarsini, and Zhong [7] is the closest to ours from a technical point of view. The focus of that paper is not entirely on the connectivity properties of best-response graphs, but it does contain the following result: in almost every game with a large number of players and exactly two actions per player, any *fixed* vertex can reach every sink via a best-response path. In Section 3.2.2 we discuss the relationship between their results and ours. Alon, Rudov, and Yariv [5] have also investigated game class sizes, showing that *dominance-solvable* games are rare when the number of actions gets large for at least one player.

In addition to our main result above on the prevalence of connected games, we also give another result on game connectivity (Proposition 3.6) which shows that, in contrast to connected games, the fraction of generic games with a pure Nash equilibrium that are acyclic vanishes super-exponentially in the number of players. *Potential* games (which are widely studied) have acyclic best-response graphs [107], so our result implies that potential games with large numbers of players are very rare. Of course, this does not suggest that potential games are somehow unimportant: they are an appropriate model for certain types of strategic interaction such as congestion [124], and many dynamics are guaranteed to converge to a pure Nash equilibrium in such games (see e.g. [83, 125]). Rather, it highlights that our results' implications for the convergence of adaptive dynamics, which we discuss below, are a consequence of the prevalence of connectedness rather than of acyclicity.

Adaptive dynamics As noted above, several influential results have established possibility and impossibility results for the convergence to equilibrium of various classes of game dynamics [13, 76, 77, 86, 154].

Establishing impossibility for a class of dynamics consists in finding collections of games such that no dynamic in the class is guaranteed to lead to equilibrium in all of them. Of course, this hinges on the parameters of the problem, namely, (i) the information that is allowed to determine players' decisions in the dynamic, (ii) the notion of convergence that is required, (iii) the type of equilibrium to which the dynamic converges, and (iv) the class of games to which the dynamic is applied.[2] With regards to (i), some of the most commonly studied classes of dynamics are those that are *uncoupled*, meaning that each player's strategy depends only on the actions of other players and on their own payoff function, or that are *completely uncoupled*, meaning that each player's strategy depends only on their own realised payoffs and actions.

There are several possibility results for dynamics that lead to mixed Nash or correlated equilibria,[3] but finding completely uncoupled, or even uncoupled, dynamics in which period-by-period play converges *almost surely* to a *pure* Nash equilibrium whenever one exists is demonstrably more challenging. The following impossibility result (see Theorem 3.40), which is due to Hart and Mas-Colell [77] and Jaggard, Lutz, Schapira, and Wright [86], is well known.[4,5]

> Impossibility for uncoupled dynamics [77, 86]: *Given a stationary, 1-recall, uncoupled dynamic for a game that has at least three players,*[6] *there is a generic*

[2]Examples of restrictions in each category include: (i) the length of a player's recall, whether the player's strategy is stationary, and whether it is uncoupled or completely uncoupled, (ii) almost sure convergence or being near an equilibrium 'most of the time', (iii) pure, mixed, correlated, or ε-Nash equilibrium, (iv) generic, dominance-solvable, potential, weakly acyclic, and so on.

[3]There are, for example, uncoupled and completely uncoupled dynamics for which the empirical distribution of play converges almost surely to the set of correlated Nash equilibria in all games [53, 58, 74, 75]. There are also uncoupled and completely uncoupled dynamics for which the behaviour probabilities converge almost surely to a mixed Nash equilibrium in all generic games [54, 59].

[4]There are stationary, 2-recall, uncoupled dynamics for which the period-by-period play converges almost surely to a pure Nash equilibrium in all games that have one [36, 77]. Jaggard, Lutz, Schapira, and Wright [86] identify other uncoupled dynamics with this convergence property in a bounded-recall synchronous setting.

[5]Babichenko [13] showed that there is no completely uncoupled dynamic for which the period-by-period play converges almost surely to a pure Nash equilibrium in every generic game that has one, or even in every large generic game that has one. On the other hand, there are completely uncoupled dynamics for which the period-by-period play is at a pure Nash equilibrium 'most of the time' in all generic games that have one [118, 155].

[6]We require that every player has at least two actions (or at least three actions if there are only three players).

game with a pure Nash equilibrium that has the same action profile space, but for which the period-by-period play under the dynamic does not almost surely converge to a pure Nash equilibrium.

Our result on connected games implies that this impossibility result does not hold when we consider 'typically' structured many-player games. In particular, we show the following (Corollary 3.43).

There exists a stationary, 1-recall, uncoupled dynamic for which the period-by-period play converges almost surely to a pure Nash equilibrium in almost every *large generic game that has one.*

This follows from combining a result of Young [153] that period-by-period play under the *best-response dynamic with inertia* (which is stationary, 1-recall, and uncoupled) converges almost surely to a pure Nash equilibrium in every weakly acyclic game with our result that almost every large generic game that has a pure Nash equilibrium is connected and hence weakly acyclic. This, of course, does not overturn the above impossibility result, but it limits its scope: the best-response dynamic with inertia is not guaranteed to converge to a pure Nash equilibrium in *every* large generic game that has one, but it does have such a guarantee in *almost every* large generic game that has one.

We remark that although we focus mostly on the best-response dynamic with inertia, since connected games are weakly acyclic it follows from our results that all existing results on the convergence of dynamics in weakly acyclic games also apply in almost every large generic game with a pure Nash equilibrium. Moreover, while these applications require only weak acyclicity, our work has recently been applied by Newton and Sawa [115] to study equilibrium selection in the context of evolutionary game theory, and in these applications genuine connectedness is required (see Section 3.6).

One possible interpretation of our results can be given by drawing a parallel with the literature on the analysis of algorithms. Algorithms are often assessed in terms of their worst-case performance. In practice, however, some algorithms such as the simplex algorithm tend to perform much better on typical problem instances than on their worst-case inputs, and this has given rise to 'average-case' or 'beyond the worst-case' analysis of algorithms [126]. Viewed in this light, the impossibility results can be interpreted as saying that uncoupled dynamics might perform very poorly on their worst-case inputs but our result shows that some such dynamics perform well in non-pathological cases. The strength of this interpretation depends to some extent on

what cases one considers non-pathological; we now explain why generic games with many players are appropriate in our setting.

Hart [72] distinguishes between three types of dynamics in games: learning, evolutionary, and adaptive. Learning requires high levels of rationality (see e.g. [88]) whereas players in evolutionary dynamics instead mechanically inherit traits (see e.g. [84, 128, 148]). Adaptive agents fall somewhere in between: they use relatively little information and take actions that respond to their environment according to simple decision heuristics in a generally improving way. Examples of adaptive dynamics include better- and best-response dynamics (which are our focus in Chapter 3), fictitious play [57], adaptive play [152], regret matching [74], regret testing [54], trial-and-error learning [155], and many more. Adaptive dynamics are most commonly and naturally applied in games with a large number of players and in which players play pure strategies. Moreover, any utility-based game for which the utility numbers are perturbed by small random shocks independently drawn from an atomless distribution will have an ordinal representation that is almost surely generic in our sense of the term. Viewed in this way, generic games with many players, which are the focus of our work, are appropriate non-pathological inputs for adaptive dynamics.[7]

Chapter 3 is based on joint work with Tom Johnston, Alex Scott, and Bassel Tarbush [87].

1.3 Induced subgraphs of induced subgraphs of large chromatic number

Given a graph with large chromatic number, what can we say about its local structure? For example, are there induced subgraphs that appear in every graph with sufficiently large chromatic number? Perhaps the first question of this type one might ask is whether every graph with large chromatic number contains a large complete subgraph. However, it has been known since the 1940s that this is not the case: Tutte [41] and Zykov [158] gave constructions of triangle-free graphs with arbitrarily large chromatic number, and Erdős [47] showed that there exist graphs with both arbitrarily large

[7] A sometimes natural condition to impose on games with many players is *anonymity*. In anonymous games, at any action profile, a player's ranking of their actions depends only on the numbers of other players selecting particular actions, not on the identities of those other players. Such games are of interest because certain interactive situations, such as congestion, can be well modelled by them, but there is no implication that games with many players outside of this class are not also of interest. Indeed, the class of weakly acyclic games, which need not be anonymous, is frequently studied in the context of adaptive dynamics with many players.

girth and chromatic number. But this raises another question: if a graph with large chromatic number does not contain a large complete subgraph, then what can we say about its induced subgraphs?

It is helpful to use the language of χ-bounded classes: a hereditary class $\mathcal{G}$ is *χ-bounded* if there is some function f such that $\chi(G) \leqslant f(\omega(G))$ for all $G \in \mathcal{G}$, where $\chi(G)$ denotes the *chromatic number* of G and $\omega(G)$ denotes the *clique number* (the maximum number of vertices in a complete subgraph of G). The study of χ-boundedness was strongly influenced by an important paper of Gyárfás [66], which put forward a number of influential conjectures. There has been a burst of recent activity in this area, and many of the conjectures have been resolved (see [133, 136]).

An important strategy for showing that a hereditary class $\mathcal{G}$ is χ-bounded has been that of 'cleaning up' its structure: given $G \in \mathcal{G}$ with very large chromatic number, we can attempt to find an induced subgraph that still has fairly large chromatic number, but is structurally simpler in some way. If this succeeds for all $G \in \mathcal{G}$, then to determine whether the class $\mathcal{G}$ is χ-bounded we need only consider these structurally simpler graphs. A very strong version of this approach was proposed independently by Galvin and Rödl (see [108]) and by Esperet (see [136]), who made the striking conjecture that every graph with large chromatic number contains either a large clique or an induced triangle-free subgraph with large chromatic number. If this were true, it would have radical consequences for the study of χ-boundedness, as it would imply that whether or not a class is χ-bounded is determined by the triangle-free graphs in the class. However, Galvin and Rödl's conjecture was recently disproved by Carbonero, Hompe, Moore, and Spirkl [35], who proved the following via a surprising new twist on a construction of Kierstead and Trotter [93].

Theorem 1.6 ([35])**.** *There are graphs of arbitrarily large chromatic number that contain neither a K_4 nor a triangle-free induced subgraph of chromatic number greater than four.*

Shortly afterwards, Briański, Davies, and Walczak [30] extended the result of Carbonero, Hompe, Moore, and Spirkl to cliques of prime order in an ingenious paper proving the following.[8]

Theorem 1.7 ([30])**.** *For every prime p, there are K_{p+1}-free graphs G of arbitrarily large chromatic number such that every K_p-free induced subgraph H of G satisfies*

[8]Building on this result, they further showed that there are classes of graphs that are χ-bounded but not polynomially χ-bounded, disproving another conjecture of Esperet [48].

$\chi(H) \leqslant \omega(H)^{\omega(H)^2}$.

In an update [31] to their paper, which appeared after the material in Chapter 4 was first made available online, Briański, Davies, and Walczak proved a strengthening of Theorem 1.7 that improves the bound on $\chi(H)$ to $\binom{3\omega(H)+1}{3}$ and removes the restriction that p is prime. This strengthening was proved independently of our work in Chapter 4.

Theorems 1.6 and 1.7 show that there are K_{p+1}-free graphs of very large chromatic number in which every induced subgraph of large chromatic number contains a copy of K_p. Thus, the 'cleaning up' strategy cannot be deployed in the straightforward way that Esperet's conjecture suggests: we cannot simply drop to an induced subgraph with large chromatic number but smaller clique number. Nevertheless, we might still hope that we can get rid of some other graphs.

In Chapter 4 we show that Theorems 1.6 and 1.7 are just the tip of the iceberg. Let us say that a graph G is *F-free* if it does not contain an induced copy of F. We will show that results such as Theorems 1.6 and 1.7 hold not just for cliques, but in fact for every non-trivial graph F.

Theorem 1.8. *For every graph F with at least one edge, there is a constant, c_F, for which there exist graphs G of arbitrarily large chromatic number and the same clique number as F such that every F-free induced subgraph of G is c_F-colourable.*

It follows that there is no general 'cleaning up' strategy to get rid of induced subgraphs: it is necessary instead to use particular properties of the class under consideration, or to focus on other structural features.[9]

Theorem 1.8 is also interesting from the perspective of Ramsey theory. A class $\mathcal{C}$ of graphs is a *vertex Ramsey class* if, for every $F \in \mathcal{C}$ and positive integer k, there is some $G \in \mathcal{C}$ such that every vertex colouring of G in k colours contains a monochromatic induced copy of F. Vertex Ramsey classes were first studied in the 1970s by Folkman [52] and Nešetřil and Rödl [112] as part of the development of a rich theory (see [109]). The general question is to determine which hereditary classes $\mathcal{C}$ of graphs are vertex Ramsey. There has been substantial progress, particularly in the case when $\mathcal{C}$ is determined by a single excluded induced subgraph, that is, when $\mathcal{C}$ is the class of F-free graphs for some F (see [52, 112, 121, 122, 129, 130]), or when $\mathcal{C}$ is defined

[9] See, for example, the discussion of diameter and Conjecture 1.10 in [135].

by excluding a clique and a tree [92, 94]. Ramsey classes of infinite graphs have also been studied in [69, 95, 96].

One of the earliest results on vertex Ramsey classes is due to Folkman [52], who proved in 1970 that the class $\mathcal{C}_r$ of K_r-free graphs is vertex Ramsey for every $r \geqslant 3$: in other words, for every K_r-free graph F and every positive integer k, there is a K_r-free graph H such that in every k-colouring of $V(H)$ some colour class contains an induced copy of F. Theorem 1.8 proves a stronger property: it shows that for every K_r-free graph F there are K_r-free graphs with arbitrarily large chromatic number in which *every* induced subgraph of large chromatic number contains an induced copy of F. This difference resembles the difference between the theorems of Van der Waerden [147] and Szemerédi [141] on monochromatic arithmetic progressions in the integers.

Interestingly, there is no edge version of Theorem 1.8: Rödl [120] proved that, for all k, every graph with sufficiently large chromatic number contains a triangle-free subgraph with chromatic number k. In particular, while the class of triangle-free graphs is both vertex and edge Ramsey [113] and satisfies the strengthened vertex Ramsey property, it does not satisfy the analogous strengthened edge Ramsey property.

If F is triangle-free, then Theorem 1.8 guarantees a triangle-free G. It is natural to ask whether one can go further and guarantee that G has the same *girth* (the length of the shortest cycle in the graph) as F. We conjecture that this is the case (see Section 4.7), but have been unable to prove it. However, in Chapter 4 we state and prove the weaker statement in which girth is replaced by *odd girth* (the length of the shortest *odd* cycle in the graph). This implies that the class $\mathcal{C}'_r$ of $\{C_3, C_5, \ldots, C_{2r+1}\}$-free graphs has the strengthened vertex Ramsey property mentioned above. We also extend and generalise Theorem 1.8 to the settings of tournaments and hypergraphs, and investigate to what extent and in what sense similar results hold for (infinite) families of graphs.

Chapter 4 is based on joint work with António Girão, Freddie Illingworth, Emil Powierski, Alex Scott, Youri Tamitegama, and Jane Tan.

1.4 Planar inducibility

The problem of determining, for a fixed small graph H, the maximum number of induced copies of H which can be contained in a graph on n vertices was first considered in 1975 by Pippenger and Golumbic [116], and has since received considerable

attention. As summarised in [49], this maximum is now known asymptotically for all graphs H on at most four vertices except the path of length 3. In the case where H is a k-cycle for $k \geqslant 5$, Pippenger and Golumbic conjectured in their paper that the maximum is asymptotically $\frac{n^k}{k^k-k}$, which they showed is an asymptotic lower bound for all k-vertex graphs H. In 2016 this conjecture was verified in the case $k = 5$ by Balogh, Hu, Lidický, and Pfender [19], and their argument was adapted by Brandt, Lidický, and Pfender (see [29]) to solve the $k = 6$ case. The current best known upper bound for general k-cycles is $2n^k/k^k$, which is due to Král', Norin, and Volec [97]. Results for other graphs H can be found in [79, 156].

A related problem on which significant progress has been made recently is to consider what happens when we restrict to planar graphs. There are two interesting problems in this setting, firstly to determine the maximum number of not necessarily induced copies of a small graph H that can be contained in a planar graph on n vertices, which we write as $f(n, H)$, and secondly to determine the same quantity when we insist that the copies of H are induced. We write $f_I(n, H)$ for this second quantity. In Chapter 5 we will present results in both settings, but with a particular focus on the latter.

The majority of work on both the induced and non-induced versions of this problem has concerned the cases where H is a path or cycle. Before detailing this progress, we highlight the early work of Alon and Caro [3] who determined $f(n, H)$ exactly for all n in the cases where H is a complete bipartite graph or a maximal planar graph with no separating triangle, and the work of Wormald [151], and independently Eppstein [45], who showed that $f(n, H) = O(n)$ if and only if H is 3-connected. In particular, Alon and Caro, and independently Wood [150], showed that $f(n, K_4) = f_I(n, K_4) = n - 3$ for all $n \geqslant 4$, and characterised the extremal graphs as those obtained from K_4 by repeatedly adding new vertices of degree 3 (while maintaining planarity).

More generally, Huynh, Joret, and Wood [85] showed that $f(n, H) = \Theta(n^{F(H)})$ for all H, where $F(H)$ denotes the so-called *flap-number* of H, and Liu [100, Corollary 6.1] showed that for all graphs H, $f_I(n, H) = \Theta(f(n, H))$. Given these two results, attention has shifted to determining $f(n, H)$ and $f_I(n, H)$ either exactly or asymptotically and significant progress has been made recently when H is a path or cycle.

Starting with cycles, in 1979 Hakimi and Schmeichel [70] showed that $f(n, C_3) = 3n - 8$ for $n \geqslant 3$ and $f(n, C_4) = \frac{1}{2}(n^2 + 3n - 22)$ for $n \geqslant 4$. Much more recently, in 2019, $f(n, C_5)$ was determined for all $n \geqslant 5$ by Győri, Paulos, Salia, Tompkins,

and Zamora [68]. Since every copy of C_3 in a graph is an induced copy, Hakimi and Schmeichel's result also determines $f_I(n, C_3)$ for all $n \geqslant 3$. We determined the values of $f_I(n, C_4)$, $f_I(n, C_5)$, and $f_I(n, C_6)$ exactly for large n in [131, 132]. In all cases the extremal graphs were determined in the respective paper, with a small correction to the extremal graphs for $f(7, C_4)$ and $f(8, C_4)$ given by Alameddine [2]. The value of $f_I(n, C_5)$ for large n was also found independently by Ghosh, Győri, Janzer, Paulos, Salia, and Zamora [61].

No further exact results are known, but a recent string of papers have yielded impressive asymptotic results for longer cycles. This progress was initiated by Cox and Martin [38], who developed a method for upper bounding $f(n, H)$ for even cycles H via a reduction to a maximum likelihood estimator problem on graphs. They went on to solve this problem for small cycles [38, 39], before Lv, Győri, He, Salia, Tompkins, and Zhu [101] solved it in general to show that $f(n, C_{2k}) = \left(\frac{n}{k}\right)^k + o(n^k)$ for all $k \geqslant 3$. Planar graphs achieving these asymptotics were first exhibited in [67]: for $m \geqslant 3$, let $F_{n,m}$ be the n-vertex graph obtained by blowing up $\lfloor \frac{m}{2} \rfloor$ pairwise non-adjacent vertices in an m-cycle to sets of as even size as possible, then $F_{n,2k}$ is a planar n-vertex graph containing $\left(\frac{n}{k}\right)^k + o(n^k)$ cycles of length $2k$. In fact, all of these $2k$-cycles are induced, so since $f_I(n, H) \leqslant f(n, H)$ for all n and H, we have also that $f_I(n, C_{2k}) = \left(\frac{n}{k}\right)^k + o(n^k)$ for all $k \geqslant 3$.

For odd cycles of length greater than 5, we can obtain a lower bound on $f_I(n, C_{2k+1})$ in the same spirit as above by observing that $F_{n,2k+1}$ is an n-vertex planar graph containing $\left(\frac{n}{k}\right)^k + o(n^k)$ induced $(2k+1)$-cycles. To obtain a better lower bound on $f(n, C_{2k+1})$ we can consider the following graphs, first defined in [60]: for each $m \geqslant 3$, let $F'_{n,m}$ be the graph obtained from $F_{n,m}$ by adding a path through each of the blown-up vertex classes. Then $F'_{n,2k}$ is an n-vertex planar graph containing $2k\left(\frac{n}{k}\right)^k + o(n^k)$ (non-induced) $(2k+1)$-cycles. Very recently, Heath, Martin, and Wells [80] extended Cox and Martin's reduction lemmas to the case of odd cycles, and showed that for $k \in \{3, 4\}$ the graphs $F'_{n,2k}$ are asymptotically optimal for the non-induced problem, that is, $f(n, C_{2k+1}) = 2k\left(\frac{n}{k}\right)^k + o(n^k)$ for $k \in \{3, 4\}$. They conjectured that this extends to all $k \geqslant 5$, and showed a general upper bound of $f(n, C_{2k+1}) \leqslant 3k\left(\frac{n}{k}\right)^k + o(n^k)$. No better upper bounds on $f_I(n, C_{2k+1})$ are known for $k \geqslant 3$.

Progress in the cases where H is a path has been restricted to the non-induced version of the problem. In what follows, we denote by P_n the path on n vertices. The case $H = P_2$ follows from Euler's formula and the case $H = P_3$ was solved by Alon and

Caro [3] as a special case of their results on complete bipartite graphs. Grzesik, Győri, Paulos, Salia, Tompkins, and Zamora [64] determined $f(n, P_4)$ exactly for all n, and Ghosh, Győri, Martin, Paulos, Salia, Xiao, and Zamora [60] determined $f(n, P_5)$ asymptotically. The latter group also conjectured the asymptotics of $f(n, H)$ for longer paths H, predicting that $f(n, P_{2k+1}) = 4k\left(\frac{n}{k}\right)^{k+1} + O(n^k)$ and $f(n, P_{2k+2}) = 8k(k+1)\left(\frac{n}{k}\right)^{k+1} + O(n^k)$ for all $k \geqslant 2$. For $k \geqslant 3$, these asymptotics are achieved by the graphs $F'_{n,2k}$, and for $k = 2$ they are achieved by a copy of P_{n-2} with the addition of two vertices of degree $n-1$.

No progress has been made on this conjecture for paths of odd length. For paths of even length, Cox and Martin [38] proved the conjecture for $H = P_7$, and for $k \geqslant 4$ gave an upper bound on $f(n, P_{2k+1})$ within a factor of around e^k of the conjectured value. This e^k factor was improved to an absolute constant by Antonir and Shapira [8] via an adapted version of Cox and Martin's method.

Our main results on $f(n, H)$ and $f_I(n, H)$ concern the cases where H is the paw graph (that is, the four-vertex graph consisting of a triangle with a pendant edge), K_4^- (that is, the graph obtained from the complete graph on four vertices by deleting one edge), or a star.

For graphs G_1 and G_2, the *graph join* of G_1 and G_2, denoted $G_1 + G_2$, is the graph consisting of a copy of G_1 and a copy of G_2 on disjoint vertex sets V_1 and V_2 respectively, where every vertex in V_1 is adjacent to every vertex in V_2. Many of the extremal graphs in our results are best expressed as graph joins. For a graph G, we write $\overline{G}$ for the *complement* of G, that is, the graph on the same vertex set as G in which vertices are neighbours if and only if they are not neighbours in G.

In Chapter 5 we show the following.

- For large n, $\overline{K_2} + C_{n-2}$ is the unique n-vertex planar graph containing the most induced paw graphs.
- For all n, we determine the n-vertex planar graphs containing the most not necessarily induced paw graphs. For large n, the unique such graph is $K_2 + P_{n-2}$.
- For all n, we determine the n-vertex planar graphs containing the most induced copies of K_4^-. For large n, the unique such graph is $K_2 + P_{n-2}$.
- For all n, we determine the n-vertex planar graphs containing the most not necessarily induced copies of K_4^-. For large n, the unique such graph is again $K_2 + P_{n-2}$.

- For fixed $k \geqslant 2$, we show that for large n the only n-vertex planar graphs containing $f_I(n, S_k)$ induced copies of S_k are $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$. For $k \in \{2, 3, 4\}$ we determine the extremal graphs exactly for *all* n.

Chapter 5 is based on unpublished work.

Chapter 2

Locally dependent percolation on $\mathbb{Z}^n$

Recall that a bond percolation model on a (possibly infinite) graph is said to be *1-independent* if for any two sets S_1 and S_2 of its edges with the property that every edge in S_1 is vertex-disjoint from every edge in S_2, it holds that the states of the edges in S_1 are independent of the states of those in S_2. As discussed in Section 1.1, such models arise naturally when renormalisation techniques are used to analyse bond percolation models in both the classical independent setting and in settings with higher-ranged dependencies.

In this chapter we will investigate for which probabilities p it holds that *every* 1-independent bond percolation model on the lattice graph $\mathbb{Z}^n$ in which each edge is open with probability p percolates. To this end, following Balister and Bollobás [16], we write $\mathcal{D}_{\geqslant p}(G)$ for the class of 1-independent bond percolation models on a (possibly infinite) graph G in which each edge is open with probability at least p, and define

$$p_{\max}(G) \coloneqq \sup\big\{p : \text{some model in } \mathcal{D}_{\geqslant p}(G) \text{ does not percolate}\big\},$$

noting that in this definition it is equivalent to ask for a model in which each edge is open with probability *exactly* p.

Balister and Bollobás [16] asked for the value of $p_{\max}(\mathbb{Z}^2)$ and, observing that $p_{\max}(\mathbb{Z}^n)$ is decreasing in n, the value of $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$. Determining either quantity exactly appears to be a very challenging problem. The best known lower bound on $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$ follows from a construction of Day, Falgas-Ravry, and Hancock [40] which gives

$$p_{\max}(\mathbb{Z}^n) \geqslant 4 - 2\sqrt{3} = 0.535898\ldots$$

for all $n \geqslant 2$, while the best rigorous lower bound on $p_{\max}(\mathbb{Z}^2)$ is due to Balister, Johnston, the author, and Scott [18] and stands at

$$p_{\max}(\mathbb{Z}^2) \geqslant \tfrac{1}{32}(35 - 3\sqrt{33}) = 0.555197\ldots$$

The authors of [18] also give a non-rigorous high confidence (p-value $< 10^{-11}$) result which suggests that, in fact, $p_{\max}(\mathbb{Z}^2) \geqslant 0.5921$.

Previously, the best upper bound on $p_{\max}(\mathbb{Z}^2)$ was 0.8639, due to Balister, Bollobás, and Walters [17], and no better upper bound on $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$ was known. The two main results of this chapter improve the upper bounds on both quantities as follows.

Theorem 1.3. $p_{\max}(\mathbb{Z}^2) \leqslant 0.8457$.

Theorem 1.4. $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n) \leqslant 0.5847$.

In Section 2.5, we build on the renormalisation argument that we use to prove Theorem 1.3 to show that under every model in $\mathcal{D}_{\geqslant 0.8457}(\mathbb{Z}^2)$ there are almost surely open cycles enclosing any bounded region of the plane. It follows that under such models the infinite component is almost surely unique and there is no infinite component in the dual graph. Moreover, we show that under every model in $\mathcal{D}_{\geqslant 0.8459}(\mathbb{Z}^2)$ the origin is in an infinite component with positive probability; we have been unable to show this all the way down to 0.8457.

The proof of Theorem 1.4 relies on the following result which translates a condition on 1-independent bond percolation models on the hypercube to one on the lattice. Here we denote by Q_n the n-dimensional hypercube graph, which has vertex set $\{0,1\}^n$ and in which vertices are joined by an edge if they differ in exactly one coordinate.

Theorem 2.1. *Let $p \in (0,1]$ and $k \in \mathbb{N}$ be constants and let $P > (1-p)^{2^{k-1}}\varphi$, where $\varphi = \frac{1}{2}(\sqrt{5}+1)$ is the golden ratio. Suppose that for any model in $\mathcal{D}_{\geqslant p}(Q_k)$, the probability that the graph is connected is at least P. Then for large enough n, every model in $\mathcal{D}_{\geqslant p}(\mathbb{Z}^n)$ percolates. In particular,*

$$\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n) \leqslant p.$$

In order to extract a concrete bound from this result we need the following lemma, which states that $p = 0.5847$ and $k = 6$ satisfy the conditions of the theorem.

Lemma 2.2. *Let φ be the golden ratio. Then there exists a constant $P > (1 - 0.5847)^{32}\varphi$ such that in any model in $\mathcal{D}_{\geqslant 0.5847}(Q_6)$, the probability that the graph is connected is at least P.*

Together, these results give Theorem 1.4. The approach we use to prove Theorem 2.1 also yields the following result about the critical probability for the asymptotically almost sure existence of a giant component in the hypercube graph Q_n under 1-independent bond percolation models, which may be of independent interest.

Theorem 2.3. *Let $p \in (0, 1]$ and $k \in \mathbb{N}$ be constants and let $P > (1-p)^{2^{k-1}}\varphi$, where φ is the golden ratio. Suppose that for any model in $\mathcal{D}_{\geqslant p}(Q_k)$, the probability that the graph is connected is at least P. Then there exists $C > 0$ such that for $n \in \mathbb{N}$, and any model in $\mathcal{D}_{\geqslant p}(Q_n)$, there is a component containing at least $C \cdot 2^n$ vertices with probability at least $1 - e^{-\Omega(n^2)}$.*

Motivated by the statement of this theorem, let

$$p_{\text{giant}} = \inf\left\{p : \exists C > 0 \colon \lim_{n\to\infty} \inf_{\mu \in \mathcal{D}_{\geqslant p}(Q_n)} \mathbb{P}_\mu(Q_n \text{ has a component of size} \geqslant C \cdot 2^n) = 1\right\}$$

be the threshold p for the asymptotically almost sure existence of a giant component in Q_n under any model in $\mathcal{D}_{\geqslant p}(Q_n)$. Then Theorem 2.3 together with Lemma 2.2 imply that $p_{\text{giant}} \leqslant 0.5847$. The best known lower bound on p_{giant} comes from a simple construction introduced by Balister and Bollobás in [16], which shows that $p_{\text{giant}} \geqslant \frac{1}{2}$.

Chapter outline After detailing some notation and preliminary results in Section 2.1, in Section 2.2 we prove Theorem 2.3 via a series of propositions and lemmas. That section contains many of the key results and ideas used in the proof of Theorem 2.1, which is dealt with in Section 2.3. The proof of Lemma 2.2 is handled in Section 2.4. We prove Theorem 1.3 in Section 2.5, and finally in Section 2.6 we discuss some possible avenues for future enquiry.

This chapter is based on joint work with Paul Balister, Tom Johnston, and Alex Scott [18].

2.1 Notation and preliminaries

Although we often work on subgraphs of Q_n, the distance between two vertices will always be the Hamming distance, or equivalently, the distance between the vertices

in Q_n. For a vertex v of a subgraph G of Q_n, define the second neighbourhood of v in G to be the set of vertices that are joined to v by a path of length exactly 2 in G, and write this as $N_G^2(v)$.

For two tuples $u \in \mathbb{Z}^i$ and $v \in \{0,1\}^j$ where $i, j \in \mathbb{N}$, we will write $u \parallel v$ to mean the concatenation of u and v; if $j = 1$, so that $v = (0)$ or $v = (1)$, then we will write $u \parallel 0$ or $u \parallel 1$ respectively. For $n \in \mathbb{N}$, we will write $[n]$ for the set $\{1, 2, \ldots, n\}$.

We now state the Chernoff bound we will use throughout this chapter (see [106] for a discussion of this result).

Lemma 2.4 (Chernoff bound)**.** *Let $n \in \mathbb{N}$, let $p \in [0,1]$, and let $X \sim \mathrm{Bin}(n,p)$. Then for all $\varepsilon \in [0,1]$,*

$$\mathbb{P}(X \leqslant (1-\varepsilon)np) \leqslant e^{-\varepsilon^2 np/2}.$$

We also use the following simple probabilistic inequality.

Lemma 2.5. *Let X be a random variable taking values in the interval $[0,N]$ for some $N \in (0,\infty)$, and suppose that $\mathbb{E}[X] \geqslant cN$ for some $c \in (0,1]$. Then for all $\varepsilon \in (0,1)$,*

$$\mathbb{P}\big(X \geqslant (1-\varepsilon)cN\big) \geqslant c\varepsilon.$$

Proof. Writing $\mathbf{1}_A$ for the indicator random variable of an event A, we have

$$\begin{aligned} cN &\leqslant \mathbb{E}[X] \\ &= \mathbb{E}[X \cdot \mathbf{1}_{\{X \geqslant (1-\varepsilon)cN\}}] + \mathbb{E}[X \cdot \mathbf{1}_{\{X < (1-\varepsilon)cN\}}] \\ &\leqslant N \cdot \mathbb{P}\big(X \geqslant (1-\varepsilon)cN\big) + (1-\varepsilon)cN. \end{aligned}$$

Rearranging yields the desired inequality. □

2.2 Giant component in Q_n

We begin in earnest by studying the existence of a giant component in the hypercube graph Q_n, and prove Theorem 2.3. A key step in the proof is a renormalisation argument which reduces the problem to one concerning a more general class of percolation models where both the vertices and edges may be open or closed, but which still maintains that the states of subgraphs induced on disjoint subsets of vertices are independent. In return for allowing vertices to be closed, edges in the renormalised models have a higher probability of being open when both endpoints are open. Formally, we will consider percolation models from the following class.

Definition 2.6. For a (possibly infinite) graph G, and for $p_v \in (0,1]$ and $p_e \in [0,1]$, define $\mathcal{D}(G, p_v, p_e)$ to be the family of percolation models on G in which

(i) vertices are open with probability at least p_v;

(ii) conditioned on their endpoints being open, edges are open with probability at least p_e; and

(iii) for all pairs of vertex-disjoint subgraphs G_1 and G_2 of G, the states of the vertices and edges in G_1 are independent of the states of the vertices and edges in G_2.

Our main result on these percolation models is the following proposition, proved in Section 2.2.2, from which Theorem 2.3 will follow almost immediately after a suitable renormalisation.

Proposition 2.7. *Let $p_v \in (0,1]$ and $p_e \in (\frac{18}{19}, 1]$ be constants. Then there exists a positive constant C such that for all $n \in \mathbb{N}$, under every model in $\mathcal{D}(Q_n, p_v, p_e)$, there is a component containing at least $C \cdot 2^n$ vertices with probability at least $1 - e^{-\Omega(n^2)}$.*

In the next section we will prove Theorem 2.3 assuming Proposition 2.7, then in Section 2.2.2 we will return to prove the proposition via a series of lemmas. In order to apply Theorem 2.3 we will make use of Lemma 2.2, which shows that the conditions of Theorem 2.3 hold when $p = 0.5847$, and which is proved in Section 2.4.

2.2.1 Proof of Theorem 2.3

We now give the proof of Theorem 2.3 assuming Proposition 2.7 holds.

Proof of Theorem 2.3 assuming Proposition 2.7. Let $n \geqslant k$ be an integer. For a given 1-independent bond percolation model μ_n on Q_n in which each edge is open with probability p, we will inductively define coupled percolation models $\mu_{n-k}, \dots, \mu_1$ on $Q_{n-k}, \dots, Q_1$ respectively, starting with μ_{n-k}. These couplings will be such that, for each $i \in \{k, \dots, n-1\}$, the existence of a component of some size s in Q_{n-i} under μ_{n-i} implies the existence of a component of size at least $s \cdot 2^i$ in Q_n under μ_n.

We say that a vertex $v \in Q_{n-k}$ is open under μ_{n-k} if the hypercube $H_v = \{v \,\|\, a : a \in \{0,1\}^k\} \subseteq Q_n$ is connected under μ_n. An edge of Q_{n-k} between two open vertices u and v is chosen to be open under μ_{n-k} if at least one of the edges between the two hypercubes H_u and H_v is open under μ_n (so that the $(k+1)$-dimensional cube $H_u \cup H_v$ is connected). For $i \in \{k, \dots, n-2\}$, define μ_{n-i-1} to be the model on

Q_{n-i-1} in which a vertex v is open if the edge between the (open) vertices $v \parallel 0$ and $v \parallel 1$ (in Q_{n-i}) is open under μ_{n-i}, and an edge uv between two open vertices is open if the edge between $u \parallel 0$ and $v \parallel 0$ or the edge between $u \parallel 1$ and $v \parallel 1$ is open. It is clear that $\mu_{n-k}, \dots, \mu_1$ satisfy the condition on component sizes from the end of the previous paragraph.

For each $i \geqslant k$, define $q_i = (1-p)^{2^i}$, so that $q_{i+1} = q_i^2$ for these i. Let $s_k = P$, and for $i \geqslant k$, let $s_{i+1} = s_i^2 - q_i$. We claim that $s_i^2 > q_i\varphi^2 > 0$ for all $i \geqslant k$, where φ is the golden ratio. Indeed, $s_k^2 = P^2 > q_k\varphi^2$ by assumption, and if $s_{i-1}^2 > q_{i-1}\varphi^2$ for $i \geqslant k+1$, then $s_i^2 = (s_{i-1}^2 - q_{i-1})^2 > q_{i-1}^2(\varphi^2-1)^2 = q_i\varphi^2$. Thus, for all $i \geqslant k$, we can define $r_i = q_i/s_i^2 \in [0,1]$.

Claim 2.8. *For all $k \leqslant i \leqslant n-1$, we have $\mu_{n-i} \in \mathcal{D}(Q_{n-i}, s_i, 1-r_i)$.*

Proof. We prove the claim by induction on i, considering first the case $i = k$. Since μ_n is a 1-independent bond percolation model, in μ_{n-k} vertices are open independently, and they are each open with probability at least $s_k = P$ by assumption. Let uv be an edge of Q_{n-k} and consider the hypercubes in Q_n corresponding to u and v. Write A for the event that all the edges in Q_n between these hypercubes are closed. Since all these edges are open independently, the probability of A occurring is q_k. Thus we have

$$\begin{aligned}\mathbb{P}(uv \text{ open } \mid u, v \text{ open}) &= \mathbb{P}(\{u, v \text{ open}\} \cap \{uv \text{ open}\})/\mathbb{P}(u, v \text{ open}) \\ &\geqslant (\mathbb{P}(u, v \text{ open}) - \mathbb{P}(A))/\mathbb{P}(u, v \text{ open}) \\ &\geqslant 1 - q_k/P^2 \\ &= 1 - r_k.\end{aligned}$$

Finally, if G_1 and G_2 are vertex-disjoint subgraphs of Q_{n-k}, then the states of the vertices and edges in each are determined by the states of the edges in two vertex-disjoint subgraphs of Q_n. Since μ_n is a 1-independent model, the states of the vertices and edges in G_1 and G_2 are independent of each other and hence $\mu_{n-k} \in \mathcal{D}(Q_{n-k}, s_k, 1-r_k)$.

Now let $k+1 \leqslant i \leqslant n-1$ and assume that $\mu_{n-i+1} \in \mathcal{D}(Q_{n-i+1}, s_{i-1}, 1-r_{i-1})$. This means each vertex in Q_{n-i+1} is open independently with probability at least s_{i-1}, and given that two endpoints of an edge of Q_{n-i+1} are open, the edge itself is open with probability at least $1 - r_{i-1}$. It follows that the probability that a vertex of Q_{n-i} is open is at least $(1-r_{i-1})s_{i-1}^2 = s_i$. If G_1 and G_2 are vertex-disjoint subgraphs

of Q_{n-i}, then the states of the vertices and edges in each are determined by the states of the vertices and edges in two vertex-disjoint subgraphs of Q_{n-i+1}. Since $\mu_{n-i+1} \in \mathcal{D}(Q_{n-i+1}, s_{i-1}, 1-r_{i-1})$, we deduce that the states of the edges and vertices in G_1 and G_2 are independent of each other.

It just remains to check that an edge uv of Q_{n-i} is present with probability at least $1-r_i$ given that u and v are open. Let $u_b = u \parallel b$ and $v_b = v \parallel b$ for each $b \in \{0,1\}$. The edge uv is open in Q_{n-i} exactly when the edges u_0u_1, v_0v_1, and at least one of u_0v_0 and u_1v_1 are open in Q_{n-i+1}. Let B be the event that u_0, u_1, v_0 and v_1 are all open in Q_{n-i+1}. Then

$$\begin{aligned}
\mathbb{P}(uv \text{ open} \mid u, v \text{ open}) &= \mathbb{P}(\{uv \text{ open}\} \cap \{u, v \text{ open}\})/\mathbb{P}(u, v \text{ open}) \\
&= \mathbb{P}(\{u_0u_1, v_0v_1 \text{ open}\} \cap \{u_0v_0 \text{ or } u_1v_1 \text{ open}\})/\mathbb{P}(u_0u_1, v_0v_1 \text{ open}) \\
&= \mathbb{P}(\{u_0u_1, v_0v_1 \text{ open}\} \cap \{u_0v_0 \text{ or } u_1v_1 \text{ open}\} \mid B)/\mathbb{P}(u_0u_1, v_0v_1 \text{ open} \mid B) \\
&\geqslant \frac{\mathbb{P}(u_0u_1, v_0v_1 \text{ open} \mid B) - \mathbb{P}(u_0v_0, u_1v_1 \text{ closed} \mid B)}{\mathbb{P}(u_0u_1, v_0v_1 \text{ open} \mid B)} \\
&\geqslant 1 - \frac{r_{i-1}^2}{(1-r_{i-1})^2} \\
&= 1 - \frac{q_{i-1}^2/s_{i-1}^4}{(s_{i-1}^2 - q_{i-1})^2/s_{i-1}^4} \\
&= 1 - r_i.
\end{aligned}$$

Hence, $\mu_{n-i} \in \mathcal{D}(Q_{n-i}, s_i, 1-r_i)$ as required. □

Claim 2.9. *We have $r_i \to 0$ as $i \to \infty$.*

Proof. For $i \geqslant k$, we have

$$\frac{1}{r_{i+1}} = \frac{s_{i+1}^2}{q_{i+1}} = \left(\frac{s_i^2 - q_i}{q_i}\right)^2 = \left(\frac{1}{r_i} - 1\right)^2.$$

If $1/r_i = \varphi^2 + \varepsilon$ for some $\varepsilon > 0$, then

$$\frac{1}{r_{i+1}} = (\varphi^2 - 1 + \varepsilon)^2 = (\varphi + \varepsilon)^2 \geqslant \varphi^2 + 2\varphi\varepsilon.$$

In particular, this can be iterated to obtain

$$\frac{1}{r_{i+j}} \geqslant \varphi^2 + (2\varphi)^j\varepsilon.$$

By assumption, $1/r_k = P^2/q_k = \varphi^2 + \varepsilon$ for some $\varepsilon > 0$ so, as $2\varphi > 1$, we have $r_i \to 0$, as claimed. □

Let I be a constant large enough that $r_I < \frac{1}{19}$. Then for all $n > I$, we have $\mu_{n-I} \in \mathcal{D}(Q_{n-I}, s_I, 1 - r_I)$ where $s_I > 0$ and $1 - r_I > \frac{18}{19}$, so by Proposition 2.7 there exists a constant C (independent of n and μ_n) such that, under μ_{n-I}, there is a component in Q_{n-I} containing at least $C \cdot 2^{n-I}$ vertices with probability at least $1 - e^{-\Omega(n^2)}$. As noted above, by the construction of $\mu_{n-k}, \ldots, \mu_{n-I}$, the existence of a component of size s in Q_{n-I} under μ_{n-I} implies the existence of a component of size at least $s \cdot 2^I$ in Q_n under μ_n. It follows that with probability at least $1 - e^{-\Omega(n^2)}$, the random subgraph of Q_n under μ_n has a component containing at least $C \cdot 2^n$ vertices, which completes the proof of Theorem 2.3 given Proposition 2.7. □

2.2.2 Proof of Proposition 2.7

In this section we will prove Proposition 2.7 via a series of lemmas. Throughout this section, when considering a model in $\mathcal{D}(Q_n, p_v, p_e)$ for some $n \in \mathbb{N}$ and $p_v, p_e \in [0, 1]$ we will write G for the associated random subgraph of Q_n, and H for the random subgraph of Q_n induced by the open vertices (thus G is the subgraph of H obtained by deleting closed edges). For any vertex v of Q_n, we will denote by H_v the subgraph of H induced by vertices with the same first coordinate as v. For each $n \in \mathbb{N}$, order the vertices of Q_n deterministically, for example using the lexicographic ordering. We will assume that subsets of $V(Q_n)$ inherit this ordering. In particular, for $v \in V(Q_n)$ we will often work with the first $\lceil p_v^2 n^2/4 \rceil$ vertices in $N^2_{H_v}(v)$ when using the fixed global ordering, and we will denote this set by $I(v)$ provided it exists. We make the following definitions.

Definition 2.10. Let $n \in \mathbb{N}$, $p_v \in (0, 1]$, $p_e \in (0, 1]$, and $\varepsilon \in (0, 1)$. Assume we are given a percolation model in $\mathcal{D}(Q_n, p_v, p_e)$.

(i) A vertex v of Q_n is *vertex-good* if it is open and $|N^2_{H_v}(v)| \geqslant p_v^2 n^2/4$. In this case we denote by $I(v)$ the first $\lceil p_v^2 n^2/4 \rceil$ vertices of $N^2_{H_v}(v)$.

(ii) A vertex v of Q_n is *ε-edge-good* if it is vertex-good and at least $(1 - \varepsilon)(2p_e - 1)p_v^2 n^2/4$ vertices of $I(v)$ are joined to v by a path of length 2 in G.

(iii) An unordered pair $\{u, v\}$ of distinct vertices of Q_n is *ε-vertex-good* if u and v are vertex-good and there are at least $(1 - \varepsilon)p_v^2 n^2/4$ vertex-disjoint paths in H, each of length at most 15, between $I(u)$ and $I(v)$.

(iv) An unordered pair $\{u, v\}$ of distinct vertices of Q_n is *ε-edge-good* if u and v are vertex-good and there are at least $(1-\varepsilon)(15p_e - 14)p_v^2 n^2/4$ vertex-disjoint paths in G between $I(u)$ and $I(v)$.

Observe that parts (ii) and (iv) of this definition are only substantive when $p_e > 1/2$ and $p_e > 14/15$ respectively, and note that whether or not a vertex v is vertex-good or a pair of vertices $\{u, v\}$ is ε-vertex-good is a property of the sites only. Also, because the definitions are based on second neighbourhoods in H_v rather than in H, whether or not v is vertex-good or ε-edge-good depends only on the sites and bonds in the copy of Q_{n-1} inside Q_n defined by vertices with the same first coordinate as v. This will become important in Section 2.3, where we will apply the upcoming lemmas to prove Theorem 1.4, our upper bound on $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$.

We will prove Proposition 2.7 by first showing that for all $\varepsilon \in (0,1)$, with probability at least $1 - e^{-\Omega(n^2)}$, every pair of distinct vertex-good vertices within distance 9 of one another form an ε-edge-good pair. We then show that there exists ε such that if this event occurs, then any two ε-edge-good vertices at distance at most 9 from one another are in the same component in G. Next, by showing that for all $\varepsilon \in (0,1)$, with probability at least $1 - e^{-\Omega(n^2)}$, every vertex in Q_n has an ε-edge-good vertex within Q_n-distance 4, we show that there exists ε such that with probability at least $1 - e^{-\Omega(n^2)}$ all ε-edge-good vertices are in the same component. To complete the proof of the proposition we show that for all $\varepsilon \in (0,1)$, with probability at least $1 - e^{-\Omega(n^2)}$ a constant fraction of the vertices in Q_n are ε-edge-good. Parts of the arguments that follow, particularly in the proofs of Lemma 2.13 and Lemma 2.17, are based on arguments of McDiarmid, Scott, and Withers in [105].

Lemma 2.11. *For all $p_v \in (0,1]$, $p_e \in (0,1]$, and $n \geqslant 12$, and all models in the class $\mathcal{D}(Q_n, p_v, p_e)$, the probability that a given vertex is vertex-good is at least $p_v^3/3$.*

Proof. Let $n \geqslant 12$ and fix a model in $\mathcal{D}(Q_n, p_v, p_e)$. Fix a vertex $v \in V(Q_n)$, condition on it being open, and denote the size of its second neighbourhood in H_v by $X_v = |N^2_{H_v}(v)|$. Let $u \in N^2_{Q_n}(v)$ have the same first coordinate as v, and let w be a common neighbour of u and v in Q_n. Since vertices of Q_n are open independently, the probability that u and w are open given that v is open is at least p_v^2. Thus, the probability that u is in $N^2_{H_v}(v)$ is at least p_v^2 and $\mathbb{E}[X_v \mid v \text{ open}] \geqslant p_v^2\binom{n-1}{2}$. Hence, applying Lemma 2.5 and noting that $p_v^2 n^2/4 \leqslant 2p_v^2\binom{n-1}{2}/3$ for $n \geqslant 12$, we have

$$\begin{aligned}\mathbb{P}\big(v \text{ vertex-good} \mid v \text{ open}\big) &\geqslant \mathbb{P}\big(X_v \geqslant 2p_v^2\tbinom{n-1}{2}/3 \mid v \text{ open}\big)\\ &\geqslant p_v^2/3.\end{aligned}$$

It follows that the probability that v is vertex-good is at least $p_v^3/3$, as required. □

Before moving on to the next lemma, we make the key observation that for any model in $\mathcal{D}(Q_n, p_v, p_e)$, if we condition on H (or, equivalently, if we condition on the set of open vertices), then the bonds follow a model in $\mathcal{D}_{\geqslant p_e}(H)$, that is, a 1-independent model on H in which edges are open with probability at least p_e. We will use this fact freely in what follows.

Lemma 2.12. *Let $p_v \in (0,1]$, $p_e \in (0,1]$, and $\varepsilon \in (0,1)$ be constants. Then there exists a constant $c = c(p_v, p_e, \varepsilon) > 0$ such that for all $n \geqslant 12$ and all models in $\mathcal{D}(Q_n, p_v, p_e)$, the probability that a given vertex is ε-edge-good is at least c.*

Proof. Let $n \geqslant 12$, fix a model in $\mathcal{D}(Q_n, p_v, p_e)$ and $v \in V(Q_n)$, and condition on the event that v is vertex-good. If $p_e \leqslant 1/2$, then the result is obvious, so assume otherwise. Let X_v be the number of elements of $I(v)$ which are joined to v by a path of length 2 in G. Given a vertex $u \in I(v)$, let w be an open neighbour of both u and v. By a union bound, the probability that at least one of the edges uw and wv is closed is at most $2(1-p_e)$, so the probability that both of them are open is at least $2p_e - 1$. Hence, $\mathbb{E}[\, X_v \mid v \text{ vertex-good}\,] \geqslant (2p_e - 1)\lceil p_v^2 n^2/4 \rceil$. Applying Lemmas 2.5 and 2.11 we have

$$\begin{aligned}\mathbb{P}(v \text{ } \varepsilon\text{-edge-good}) &= \mathbb{P}(X_v \geqslant (1-\varepsilon)(2p_e-1)p_v^2 n^2/4 \mid v \text{ vertex-good}) \\ &\quad \cdot \mathbb{P}(v \text{ vertex-good}) \\ &\geqslant \varepsilon(2p_e - 1) \cdot p_v^3/3,\end{aligned}$$

which completes the proof of the lemma. □

Lemma 2.13. *Let $p_v \in (0,1]$, $p_e \in (0,1]$, and $\varepsilon \in (0,1)$ be constants. Then, in any model in $\mathcal{D}(Q_n, p_v, p_e)$, with probability at least $1 - e^{-\Omega(n^2)}$ all pairs $\{u, v\}$ of distinct, vertex-good vertices of Q_n at distance at most 9 from one another are ε-vertex-good.*

Proof. Note that we may assume n is large. Fix a model in $\mathcal{D}(Q_n, p_v, p_e)$, and let u and v be distinct vertices of Q_n at distance at most 9 from one another. We will show that the probability that u and v are both vertex-good, but the pair $\{u, v\}$ is not ε-vertex-good, is at most $e^{-\Omega(n^2)}$. Since there are at most 4^n pairs of vertices in Q_n, a union bound will complete the proof.

Roughly, we will choose large sets $A \subseteq I(u)$ and $B \subseteq I(v)$, pair up the vertices in A and B, and then to each pair (a, b) associate many paths from a to b in Q_n. The paths for a given pair will be vertex-disjoint except at their endpoints so the internal

vertices will all be open independently. This means it is highly likely that for at least one of these paths all the vertices will be open. The paths for different pairs will be vertex-disjoint and so there will be a vertex-open path for each pair independently. Since A and B are large, a Chernoff bound will complete the proof.

We identify Q_n with $\mathcal{P}(n)$ in the natural way and assume without loss of generality that $u = \emptyset$ and $v = [d]$, where $d \leqslant 9$ is the distance between u and v. Start by revealing the states of all vertices at distance (in Q_n) at most 2 from u or v. Assume that both vertices are vertex-good, so that we have large sets $I(u)$ and $I(v)$ of vertices that are at distance 2 in H from u and v respectively. Let $N = N(p_v, \varepsilon)$ be a large constant, and construct sets $A \subseteq I(u)$ and $B \subseteq I(v)$ as follows.

1. Remove from $I(u)$ or $I(v)$ any sets which differ from u or v in any of the elements $1, \ldots, 13N + 9$. Let the new sets be A_0 and B_0 respectively.
2. Arbitrarily delete elements from the larger of A_0 and B_0 until they are of equal size to obtain A and B.

We first show that A and B are large. As elements of $I(u)$ or $I(v)$ differ from u and v in exactly two elements, the first step above removes at most $(13N + 9)n$ elements from each of $I(u)$ and $I(v)$. As $|I(u)| = |I(v)| = \lceil p_v^2 n^2/4 \rceil$, we have $|A| = |B| = \min\{|A_0|, |B_0|\} \geqslant (1 - \varepsilon/2)p_v^2 n^2/4$ for large enough n.

Arbitrarily pair up elements of A and B to obtain $(a_1, v \cup b_1), \ldots, (a_\ell, v \cup b_\ell)$ where $\ell = |A| = |B|$ and $a_i \in A$ and $v \cup b_i \in B$ for $i = 1, \ldots, \ell$. Here each a_i and b_i is a pair of elements in $\{13N + 10, \ldots, n\}$. To each pair $(a, b) = (a_i, b_i)$ we associate an integer in [13] which we will use to construct the paths associated to the pair. Given a pair (a, b), label the elements of a as α_1 and α_2, and the elements of b as β_1 and β_2, where we assume that $\alpha_1 < \alpha_2$ and $\beta_1 < \beta_2$. If $\alpha_1, \alpha_2, \beta_1$, and β_2 are all distinct, then they have six possible orderings when they are sorted into ascending order, and we associate each order with a distinct number in [6] arbitrarily. If $\alpha_1 = \beta_1$ but $\alpha_2 \neq \beta_2$, then there are two possible orderings and we associate them to 7 and 8. Continuing in this manner, we can associate a unique integer in [13] to each of the 13 possible orderings.

Given a pair (a, b) and $j \in [N]$, define a path $P_{a,j}$ from a to $v \cup b$ as follows. Let $k \in [13]$ be the integer associated to (a, b) as above, and let $x = (k - 1)N + j + 9$. Begin with the vertex a followed by $a \cup \{x\}$, $a \cup \{\beta_1, x\}$, and $a \cup b \cup \{x\}$, ignoring duplicate vertices. From here we add the elements of $v = [d]$ one by one in increasing order until we reach $a \cup b \cup v \cup \{x\}$. Finally, we add the vertices $b \cup v \cup \{\alpha_2, x\}$,

$b \cup v \cup \{x\}$, and $b \cup v$. Note that the path contains at most 16 vertices, with the exact number depending on $|a \cap b|$ and d.

Claim 2.14. *Let $a, a' \in A$ and $j, j' \in [N]$ be such that $a \neq a'$. Then the paths $P_{a,j}$ and $P_{a',j'}$ are vertex-disjoint, and if $j \neq j'$, then the paths $P_{a,j}$ and $P_{a,j'}$ are vertex-disjoint except at their endpoints.*

Proof. Let $v \cup b \in B$ be the vertex paired with a, and let $k \in [13]$ be the integer associated to the pair (a, b). We start with the second part of the claim. Every vertex in the path $P_{a,j}$ except the two end vertices contains a unique element $x = (k-1)N + j + 9$ in $[10, 13N+9]$, and so $P_{a,j}$ and $P_{a,j'}$ can only share an internal vertex if $(k-1)N + j + 9 = (k-1)N + j' + 9$, which implies $j = j'$. Moreover, it cannot be the case that an internal vertex of $P_{a,j}$ is an endpoint of $P_{a,j'}$ as the endpoints do not contain any element of $[10, 13N + 9]$.

Turning to the first part of the claim, let z be a vertex on the path $P_{a,j}$. We will show that the set a is uniquely determined. First, if $z \cap [10, 13N + 9] = \emptyset$, then $z = a$ or $z = b \cup v$, and a is clearly uniquely determined in both cases. Thus, suppose that z contains an element $x \in [10, 13N + 9]$. Since $x > 9$, we have $x \notin v$ and since $x \leqslant 13N + 9$, we have $x \notin a \cup b$. Hence, $x = (k-1)N + j + 9$ and we can read off the value of k. If $z \cap v = \emptyset$, then z is one of $a \cup \{x\}$, $a \cup \{\beta_1, x\}$, or $a \cup b \cup \{x\}$. Using the value of k and the size of z we can deduce which case we are in, and further, which of the elements form the pair a. Similarly, if $v \subseteq z$, then z is one of $a \cup b \cup v \cup \{x\}$, $b \cup v \cup \{\alpha_2, x\}$, or $b \cup v \cup \{x\}$, and using the value of k and the size of z we can deduce which case we are in, and further, which of the elements form the pair b (which then determines a). Finally, if $\emptyset \neq z \cap v \neq v$ then $z \setminus v = a \cup b \cup \{x\}$ and again we can determine where we are in the path (from $|z \cap v|$) and the value of a from k. □

For each $a \in A$, let C_a be the event that all the vertices are open in at least one of the paths $P_{a,j}$, $j \in [N]$. By construction, every vertex on the paths except for the endpoints are at distance at least 3 from u and v, and they are still open independently with probability at least p_v. The paths contain at most 14 internal vertices and hence each path is open with probability at least p_v^{14}. As the paths are disjoint except at their endpoints, the vertices in each path are open independently and, by choosing N large enough relative to ε and p_v, the event C_a has probability at least $1 - \varepsilon/2$ for all a.

By Claim 2.14, for distinct $a, a' \in A$ the events C_a and $C_{a'}$ depend on disjoint sets of vertices and are independent. By a Chernoff bound, C_a holds for at least $(1-\varepsilon)p_v^2 n^2/4$ values of a with probability at least $1 - e^{-\Omega(n^2)}$. There are at most 4^n pairs of vertices in Q_n, so by a union bound, every pair $\{u, v\}$ of distinct, vertex-good vertices in Q_n at distance at most 9 from one another is ε-vertex-good with probability at least $1 - e^{-\Omega(n^2)}$, as required. This completes the proof of Lemma 2.13. □

Lemma 2.15. *Let $p_v \in (0, 1]$, $p_e \in (\frac{14}{15}, 1]$, and $\varepsilon \in (0, 1)$ be constants. Then, under every model in $\mathcal{D}(Q_n, p_v, p_e)$, with probability at least $1 - e^{-\Omega(n^2)}$ every pair of distinct, vertex-good vertices of Q_n at distance at most 9 from one another form an ε-edge-good pair.*

Proof. Let $n \in \mathbb{N}$ and fix a model in $\mathcal{D}(Q_n, p_v, p_e)$. Reveal the graph H (i.e. reveal the states of the vertices but not the states of the edges) and suppose that every pair of distinct vertex-good vertices at distance at most 9 from one another form an $(\varepsilon/2)$-vertex-good pair. By Lemma 2.13, the probability that this does not hold is at most $e^{-\Omega(n^2)}$. Let u and v be distinct vertex-good vertices in H at distance at most 9 from one another in Q_n. We will show that the probability that the pair $\{u, v\}$ is not ε-edge-good (given the graph H) is at most $e^{-\Omega(n^2)}$, then since there are at most 4^n choices for u and v, the probability that there are two distinct vertex-good vertices which do not form an ε-edge-good pair is also at most $e^{-\Omega(n^2)}$.

By our assumption on H, the pair $\{u, v\}$ is an $(\varepsilon/2)$-vertex-good pair so there is a set of $\lceil (1 - \varepsilon/2)p_v^2 n^2/4 \rceil$ vertex-disjoint paths in H, each of length at most 15, between $I(u)$ and $I(v)$. Since the paths are vertex-disjoint and we have only conditioned on the realisation of H, it follows from the fact that the model is in $\mathcal{D}(Q_n, p_v, p_e)$ that each path is open independently. Further, the probability that a given path is open is at least $15p_e - 14$, and applying a Chernoff bound shows that $\{u, v\}$ is an ε-edge-good pair with probability at least $1 - e^{-\Omega(n^2)}$, as required. □

In the proof of the next lemma we will make use of the following theorem of Wilson which gives the maximum number of edge-disjoint copies of K_4 that can be packed into K_n (see also [33]).

Theorem 2.16 ([149])**.** *If n is sufficiently large, then the maximum cardinality of a set of 4-subsets of $[n]$ which pairwise intersect in at most one element is $\left\lfloor \frac{n}{4} \left\lfloor \frac{n-1}{3} \right\rfloor \right\rfloor - 1$ if $n \equiv 7, 10 \bmod 12$ and $\left\lfloor \frac{n}{4} \left\lfloor \frac{n-1}{3} \right\rfloor \right\rfloor$ otherwise.*

In fact, it will be sufficient for our purposes that this maximum cardinality is $\Omega(n^2)$, and it is not difficult to construct suitable sets of this size.

Lemma 2.17. *Let $p_v \in (0,1]$ and $p_e \in (\frac{18}{19},1]$ be constants. Then there exists a constant $\varepsilon \in (0,1)$ such that, under every model in $\mathcal{D}(Q_n, p_e, p_v)$, the probability that all ε-edge-good vertices in Q_n are in the same component in G is at least $1-e^{-\Omega(n^2)}$.*

Proof. Let $\varepsilon \in (0,1)$ be a constant small enough that $(1-\varepsilon)(19p_e-16) > 2$. Let $n \in \mathbb{N}$ and fix a model in $\mathcal{D}(Q_n, p_v, p_e)$. Note that we may assume n is large in terms of p_v, p_e, and ε. We will start by showing that if u and v are distinct vertices of Q_n at distance at most 9 from one another such that u and v are ε-edge-good vertices and $\{u,v\}$ is an ε-edge-good pair, then u and v are in the same component in G.

Indeed, let u and v be such vertices, and note that they are also both vertex-good. As $\{u,v\}$ is ε-edge-good, there is a set S of at least $(1-\varepsilon)(15p_e-14)p_v^2n^2/4$ vertex-disjoint paths from $I(u)$ to $I(v)$ in G. Let T_1 be the set of pairs (a,b) such that a and b are the endpoints of a path in S, where $a \in I(u)$ and $b \in I(v)$. Arbitrarily pair up the remaining elements of $I(u)$ and $I(v)$ and let the set of all these pairs be T, so that $|T| = \lceil p_v^2n^2/4 \rceil$.

Let T_2 be the set of pairs in T whose first entry has a path to u in G of length 2, and let T_3 be the set of pairs whose second entry has a path to v in G of length 2. Then, since u and v are ε-edge-good vertices, we have $|T_2|, |T_3| \geqslant (1-\varepsilon)(2p_e-1)p_v^2n^2/4$. Thus,

$$\begin{aligned} |T_1|+|T_2|+|T_3| &\geqslant (1-\varepsilon)[2(2p_e-1)+(15p_e-14)]p_v^2n^2/4 \\ &= (1-\varepsilon)(19p_e-16)p_v^2n^2/4. \end{aligned}$$

By our choice of ε, we have $(1-\varepsilon)(19p_e-16) > 2$, so if n is large enough in terms of p_v, p_e, and ε, then there exists a pair $(a,b) \in T_1 \cap T_2 \cap T_3$. It follows that u and v are in the same component in G since there is a walk from one to the other via this a and b.

We now prove the following claim.

Claim 2.18. *With probability at least $1-e^{-\Omega(n^2)}$, every vertex in Q_n has an ε-edge-good vertex within distance 4.*

Proof. First note that any set of vertices in Q_n which are pairwise at distance at least 5 from one another are ε-edge-good independently. Fix a vertex v in Q_n and

assume without loss of generality that $v = \mathbf{0}$. By Theorem 2.16 there exists a set S of vertices at distance 4 from v in Q_n which are pairwise at distance at least 6 in Q_n with $|S| = \Omega(n^2)$. The vertices in S are ε-edge-good independently, and by Lemma 2.12 (noting that we may assume that n is large), there exists a constant $c > 0$ such that each of them is ε-edge-good with probability at least c. Hence, the probability that at least one of the vertices in S is ε-edge-good is at least $1 - e^{-\Omega(n^2)}$, from which the claim follows by a union bound. □

Now suppose that every pair of distinct vertex-good vertices of Q_n at distance at most 9 from one another form an ε-edge-good pair and every vertex in Q_n has an ε-edge-good vertex within distance 4. We will show that under these assumptions, if u and v are distinct ε-edge-good vertices of Q_n, then they are in the same component of G. Since our assumptions hold with probability $1 - e^{-\Omega(n^2)}$, this will complete the proof.

Fix a path between u and v in Q_n, say $x_0x_1 \dots x_k$ where $x_0 = u$ and $x_k = v$. Each of $x_5, \dots, x_{k-5}$ are within distance 4 of ε-edge-good vertices, say $y_5, \dots, y_{k-5}$ respectively. The vertices $x_0, y_5, y_6, \dots, y_{k-5}, x_k$ are all vertex-good, and since each vertex is within distance 9 of the previous one, each pair of consecutive vertices is ε-edge-good by assumption. By the first part of the proof, the vertices in each such pair are therefore in the same component in G. It follows that u and v are in the same component and we have completed the proof of Lemma 2.17. □

We are now ready to prove Proposition 2.7.

Proof of Proposition 2.7. By Lemma 2.17 it is sufficient to show that for all constants $\varepsilon \in (0, 1)$ there exists a positive constant C such that for all $n \in \mathbb{N}$ and models in $\mathcal{D}(Q_n, p_e, p_v)$, there are at least $C \cdot 2^n$ ε-edge-good vertices in Q_n with probability at least $1 - e^{-\Omega(n^2)}$. Fix an $\varepsilon \in (0, 1)$, let $n \geqslant 6$, and consider a model in $\mathcal{D}(Q_n, p_e, p_v)$.

We observed in the proof of Claim 2.18 that any set of vertices which are pairwise at distance at least 5 in Q_n are ε-edge-good independently of one another, and by Lemma 2.12 (noting that we may assume that n is large) there exists a constant $c > 0$ such that the probability that a given vertex of Q_n is ε-edge-good is at least c. Consider the fifth power $Q_n^{(5)}$ of Q_n, that is the graph with vertex set $V(Q_n)$ and edges between vertices which are at distance at most 5 from one another in Q_n. This graph is Δ-regular for some $\Delta = O(n^5)$, so by Brooks' theorem [32] we can properly vertex-colour $Q_n^{(5)}$ with Δ colours.

Each colour class consists of vertices which are pairwise at distance at least 5 from one another in Q_n, so they are ε-edge-good independently of one another. Hence, by a Chernoff bound, there exists a constant $C' = C'(p_v, p_e, \varepsilon) > 0$ such that if D is a colour class with $|D| \geqslant n^2$, then at least $C'|D|$ of the vertices in D are ε-edge-good with probability at least $1 - e^{-\Omega(n^2)}$. Since $\Delta = O(n^5)$ it follows by a union bound that, with probability at least $1 - e^{-\Omega(n^2)}$, in every colour class of size at least n^2 at least a C' proportion of the vertices are ε-edge-good. The total number of vertices in colour classes of size less than n^2 is at most $\Delta \cdot n^2 = O(n^7)$, so there exists a constant $C > 0$ such that with probability at least $1 - e^{-\Omega(n^2)}$, at least a C proportion of the vertices of Q_n are ε-edge-good, as required to prove Proposition 2.7. □

With the proof of Proposition 2.7 now complete, this concludes the proof of Theorem 2.3.

2.3 Percolation in $\mathbb{Z}^n$

In this section we build on the methods and results of Section 2.2 to prove Theorem 2.1. As was the case with Theorem 2.3, our proof of Theorem 2.1 rests on a renormalisation argument which reduces the problem to one concerning percolation models of the form given in Definition 2.6. We now state the analogue of Proposition 2.7 for this setting, before proceeding to prove Theorem 2.1 from the proposition by mimicking the proof of Theorem 2.3 from Proposition 2.7. Following this, we will prove the proposition.

Proposition 2.19. *Let $p_v \in (0, 1]$ and $p_e \in (\frac{18}{19}, 1]$ be constants. Then, for large enough n, all models in $\mathcal{D}(\mathbb{Z}^2 \times \{0,1\}^{n-2}, p_v, p_e)$ percolate.*

Proof of Theorem 2.1 assuming Proposition 2.19. To prove the theorem it is sufficient to show that if n is large, then every 1-independent bond percolation model on $\mathbb{Z}^2 \times \{0,1\}^{n-2}$ in which each edge is open with probability p percolates. Let n be large and let μ_n be such a model on $\mathbb{Z}^2 \times \{0,1\}^{n-2}$. Define a percolation model μ_{n-k} on $\mathbb{Z}^2 \times \{0,1\}^{n-k-2}$ by defining a vertex v to be open if the hypercube $H_v = \{v \parallel a : a \in \{0,1\}^k\}$ is connected under μ_n, and defining an edge of $\mathbb{Z}^2 \times \{0,1\}^{n-k-2}$ between two open vertices u and v to be open if at least one of the 2^k edges in $\mathbb{Z}^2 \times \{0,1\}^{n-2}$ between H_u and H_v is open under μ_n. For $i \in \{k, \dots, n-3\}$, recursively define a model μ_{n-i-1} on $\mathbb{Z}^2 \times \{0,1\}^{n-i-3}$ by defining a vertex v to be open if the edge between $v \parallel 0$ and $v \parallel 1$ is open in $\mathbb{Z}^2 \times \{0,1\}^{n-i-2}$ under μ_{n-i}, and setting

an edge uv to be open if the edge between $u \parallel 0$ and $v \parallel 0$ or the edge between $u \parallel 1$ and $v \parallel 1$ is open in $\mathbb{Z}^2 \times \{0,1\}^{n-i-2}$ under μ_{n-i}.

Let q, q_i, s_i, and r_i for $i \geqslant k$ be as in the proof of Theorem 2.3. Then we again have that $r_i \to 0$ as $i \to \infty$ and, for all $k \leqslant i \leqslant n-2$, we have $s_i > 0$ and $\mu_{n-i} \in \mathcal{D}(\mathbb{Z}^2 \times \{0,1\}^{n-i-2}, s_i, 1-r_i)$. Thus, we can take I to be a constant large enough that $r_I < \frac{1}{19}$, so that for $n \geqslant I+2$ we have $\mu_{n-I} \in \mathcal{D}(\mathbb{Z}^2 \times \{0,1\}^{n-I-2}, s_I, 1-r_I)$ where $s_I > 0$ and $1 - r_I > \frac{18}{19}$, and Proposition 2.19 applies. Hence, μ_{n-I} percolates. By the construction of μ_{n-I} from μ_n, the existence of an infinite component in $\mathbb{Z}^2 \times \{0,1\}^{n-I-2}$ under μ_{n-I}, implies the existence of an infinite component in $\mathbb{Z}^2 \times \{0,1\}^{n-2}$ under μ_n, and hence μ_n percolates, which completes the proof of the theorem. □

We now turn our attention to Proposition 2.19.

Proof of Proposition 2.19. To simplify the notation, we will prove the equivalent statement that for large enough n, all models in $\mathcal{D}(\mathbb{Z}^2 \times \{0,1\}^{n-1}, p_v, p_e)$ percolate. To begin, apply Lemma 2.17 to find $\varepsilon \in (0,1)$ such that for each $n \in \mathbb{N}$, under every model in $\mathcal{D}(Q_n, p_e, p_v)$, with probability at least $1 - e^{-\Omega(n^2)}$ all ε-edge-good vertices in Q_n are in the same component. Now let n be large and consider $\mu \in \mathcal{D}(\mathbb{Z}^2 \times \{0,1\}^{n-1}, p_v, p_e)$. Let G be the random subgraph of $\mathbb{Z}^2 \times \{0,1\}^{n-1}$ associated with μ. To each edge $e \in E(\mathbb{Z}^2)$, we associate the natural copy of Q_n in $\mathbb{Z}^2 \times \{0,1\}^{n-1}$, which we denote by C_e. The subgraph of G induced on $V(C_e)$ follows a model in $\mathcal{D}(Q_n, p_v, p_e)$, so we can define a bond percolation model $\nu = \nu(\mu)$ on $\mathbb{Z}^2$ coupled to μ by declaring $e \in E(\mathbb{Z}^2)$ to be open if the following two conditions hold:

(a) in the subgraph of G induced on $V(C_e)$, all ε-edge-good vertices of C_e are in the same component; and

(b) if we partition $V(C_e)$ into two equal parts according to the first two coordinates of the vertices, then there are ε-edge-good vertices of C_e in both halves.

Since $\mu \in \mathcal{D}(\mathbb{Z}^2 \times \{0,1\}^{n-1}, p_v, p_e)$, it is clear that ν is a 1-independent bond percolation model on $\mathbb{Z}^2$. By our choice of ε, for each edge e condition (a) holds with probability at least $1 - e^{-\Omega(n^2)}$. For condition (b), recall that sets of vertices of Q_n which are pairwise at distance at least 5 are ε-edge-good independently. It is straightforward to see that there exists such a subset of $V(C_e)$ in which $\Omega(n^2)$ vertices are taken from each half of the cube, so applying Lemma 2.12 and using a Chernoff bound on each half we find that for each edge e condition (b) holds with probability at least $1 - e^{-\Omega(n^2)}$. Thus ν is a 1-independent bond percolation model on $\mathbb{Z}^2$ in which each

edge is open with probability at least $1 - e^{-\Omega(n^2)}$. Since $p_{\max}(\mathbb{Z}^2) < 1$, for example by Theorem 1.3 or as shown in [17] and in [99], we see that ν percolates if n is large enough in terms of p_v and p_e.

It remains to show that percolation of ν implies percolation of μ. We claim that for any collection of connected, open edges under ν, the ε-edge-good vertices in the associated hypercubes are in the same component under μ. An infinite component under ν contains infinitely many disjoint edges and the hypercubes of each of these must contain two ε-edge-good vertices, so this immediately implies that μ percolates when ν percolates. To prove the claim, suppose that $v \in V(\mathbb{Z}^2 \times \{0,1\}^{n-1})$ is ε-edge-good when considered as a vertex in C_e for some $e \in E(\mathbb{Z}^2)$. Let f be another edge of $\mathbb{Z}^2$ such that $v \in V(C_f)$ (so that e and f have a common vertex in $\mathbb{Z}^2$). Since the definition of a vertex in a copy of Q_n being ε-edge-good depends only on the Q_{n-1} subgraph with the same first coordinate as that vertex, it follows that v is also ε-edge-good when considered as a vertex of C_f. Now if e and f are distinct open edges of $\mathbb{Z}^2$ with a common vertex, then by condition (b) applied to e or f there exists an ε-edge-good vertex in $V(C_e) \cap V(C_f)$. By condition (a) applied to e and f, all the ε-edge-good vertices in $V(C_e) \cup V(C_f)$ are in the same component in G, and the claim follows immediately. □

2.4 Extracting a concrete bound

In this section we prove Lemma 2.2, which allows us to extract concrete bounds from Theorems 2.1 and 2.3.

Proof of Lemma 2.2. Given a model in $\mathcal{D}_{\geqslant 0.5847}(Q_6)$, we start by constructing a model in $\mathcal{D}(Q_3, p_v, p_e)$ for some $p_v, p_e \in (0,1]$ by applying the first step (with $k = 3$) of the renormalisation process used in the proof of Theorem 2.3. That is, we declare a vertex v of Q_3 to be open if the cube $\{v \parallel a : a \in \{0,1\}^3\}$ is connected under the model on Q_6, and an edge of Q_3 between two open vertices to be open if at least one of the eight edges between the two cubes corresponding to its endpoints is open. If $P_0 > 0$ is a constant such that for any model in $\mathcal{D}_{\geqslant 0.5847}(Q_3)$ the probability that the associated random subgraph is connected is at least P_0, then by Claim 2.8 the renormalised model on Q_3 is in $\mathcal{D}(Q_3, P_0, 1 - 0.4153^8/P_0^2)$, where $0.4153 = 1 - 0.5847$.

It is clear that this renormalisation has the property that if the random subgraph associated to the renormalised model forms a connected spanning subgraph of Q_3, then the random subgraph of Q_6 associated to the original model is connected. Moreover,

if we condition on all sites being open in the renormalised model, then the bonds follow a model in $\mathcal{D}_{\geqslant 1-0.4153^8/P_0^2}(Q_3)$. Thus if $P_1 > 0$ is a constant such that for any model in this class, the probability that the random subgraph of Q_3 is connected is at least P_1, then the probability that Q_6 is connected under the original model is at least $P_0^8 P_1$.

In order to obtain suitable values for P_0 and P_1, we construct a linear program that is satisfied by any 1-independent model on Q_3 (with edge probability exactly p) by removing the non-linear conditions. More specifically, for each subset S of the edges of Q_3, let x_S denote the probability that the set of open edges is exactly S, and y_S denote the probability that S is a subset of the set of open edges. Let $\mathcal{C}$ be the collection of all subsets S of $E(Q_3)$ that form a connected spanning subgraph of Q_3. We consider the following linear programming problem.

$$\begin{aligned} \text{Minimise:} \quad & \textstyle\sum_{S\in\mathcal{C}} x_S \\ \text{Subject to:} \quad & x_S \geqslant 0, \\ & y_S = \textstyle\sum_{T\supseteq S} x_T, \\ & y_{S\cup\{e\}} = p \cdot y_S, \\ & y_\emptyset = 1, \end{aligned}$$

where S runs over all subsets of $E(Q_3)$ and e runs over all edges that are vertex-disjoint from all edges in S.

It is clear that the above conditions hold in any 1-independent model with edge probability p. Indeed, the only conditions missing are the non-linear constraints $y_{S\cup T} = y_S y_T$ when S and T are sets of edges sharing no vertex and $|S|, |T| \geqslant 2$. Thus the solution to the linear programming problem gives a lower bound on the minimum probability that the open edges in any 1-independent bond percolation model on Q_3 in which edges are open with probability p form a connected spanning subgraph of Q_3. As the existence of a spanning connected open subgraph is an increasing event, this bound also holds when the edges are open with probability at least p as we can independently delete edges so as to ensure edges are open with probability exactly p.

Running the above LP problem using the `Gurobi` optimisation package and $p = 0.5847$ gives a lower bound[1] on P_0 of 0.0463. Running the LP problem again with $p = 0.5872 < 1 - 0.4153^8/0.0463^2$ gives a lower bound on P_1 of 0.0497. Hence we may take $P = 10^{-12} < 0.0497 \cdot (0.0463)^8$ as a lower bound on the probability of Q_6 being

[1]We also ran the dual programs and checked that the dual solutions provided by `Gurobi` were feasible.

connected under any model in $\mathcal{D}_{\geqslant 0.5847}(Q_6)$, and $P > 9.93 \times 10^{-13} > 0.4153^{32} \cdot \varphi$, as required. □

2.5 Upper bounds on $p_{\max}(\mathbb{Z}^2)$

We now prove Theorem 1.3 which gives an improved upper bound on $p_{\max}(\mathbb{Z}^2)$. The previous best known bound follows from a 3×3 renormalisation argument due to Balister, Bollobás, and Walters [17]. We follow their renormalisation idea, using a 2×2 renormalisation, but introduce the idea of solving a linear program to lower bound the probability that edges in the renormalised model are present. Our renormalisation is also more sophisticated, utilising an additional random element independent of the original model and giving rise to a renormalised model in which edges are open with varying probabilities.

Proof of Theorem 1.3. Tile $\mathbb{Z}^2$ with 2×2 squares $S_{(i,j)} = \{2i, 2i+1\}\times\{2j, 2j+1\}$ and form a new 1-independent model on $\mathbb{Z}^2$ by identifying $S_{(i,j)}$ with the site $(i, j) \in \mathbb{Z}^2$ and joining neighbouring sites $u, v \in \mathbb{Z}^2$ if a certain 'good event' E_{uv} occurs in the corresponding 4×2 or 2×4 rectangle $S_u \cup S_v$ in the original 1-independent model. The good event E_{uv} will be defined to be the existence of an open component in the rectangle $S_u \cup S_v$ that has a 'large' intersection with both S_u and S_v. As long as any two 'large' subsets of any S_u intersect, these events will have the property that an open path from some v to some w in the renormalised grid will imply the existence of an open path from some site in S_v to some site in S_w in the original model.

Clearly, containing at least three out of four sites in S_u would be a sufficient condition to be 'large', but we can do better. Fix for each $u \in \mathbb{Z}^2$ a subset $\mathcal{L}_u \subseteq S_u$ of size 3, i.e. one of the subsets shown in Figure 2.1. We define a large subset of S_u to be one that intersects $\mathcal{L}_u$ in at least two sites. Note that this includes any set containing three or more sites in S_u, but certain 2-element subsets of S_u are now also deemed large and, of course, any two large subsets of S_u intersect.

We now come to the choice of $\mathcal{L}_u$. These could be chosen to be all (translates of) the same fixed set, but we can improve the bound slightly by having $\mathcal{L}_u$ vary, and indeed vary randomly. Hence, we shall choose $\mathcal{L}_u$ randomly, independently for each u, and also independently of the state of the original model. We note that under these assumptions, the renormalised model is still 1-independent as vertex-disjoint sets of renormalised edges depend on vertex-disjoint subgraphs of the original model, and on the choice of $\mathcal{L}_u$ for disjoint sets of sites u.

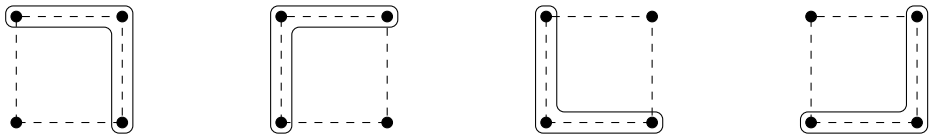


Figure 2.1: The sets $S_u^{r,s}$

Even the probability distributions for the choice of $\mathcal{L}_u$ will vary, depending on u, so our renormalised bonds will not be open with the same probabilities. Hence, to inductively renormalise we will need to consider 1-independent models with differing edge probabilities for each edge. However, it will be enough for our purposes to consider models with at most two distinct edge probabilities. More precisely, define $\mathcal{D}_{p,p'}$ as the set of 1-independent bond percolation models on $\mathbb{Z}^2$ in which each edge uv is open with probability p if u and v both lie in the same 2×2 square S_w, and with probability p' otherwise (see Figure 2.2). Similarly we define $\mathcal{D}_{\geqslant p,\geqslant p'}$ as the set of 1-independent models where the edge probabilities are at least p and p' respectively.

We now give the probability distribution for the $\mathcal{L}_u$. For $r, s \in \{0, 1\}$ let

$$S_{(i,j)}^{r,s} = S_{(i,j)} \setminus \{(2i + r, 2j + s)\},$$

as shown in Figure 2.1. Fix $\theta \in [0, 1]$, then for $u \in \mathbb{Z}^2$ with $u \equiv (r, s) \bmod 2$, define $\mathcal{L}_u$ so that

$$\mathcal{L}_u = \begin{cases} S_u^{r,s} & \text{with probability } 1 - \theta, \\ S_u^{1-r,1-s} & \text{with probability } \theta. \end{cases}$$

Figure 2.2 shows the situation when the first case always holds (i.e., when $\theta = 0$). An equivalent definition is that we initially set $\mathcal{L}_u = S_u^{r,s}$, and then independently with probability θ rotate each $\mathcal{L}_u$ by 180°.

Consider the 4×2 rectangle $R = \{0, 1, 2, 3\} \times \{0, 1\}$ and let E be the set of ten edges it induces. Let $\mathcal{C}_0$ be the set of subsets of E which give an open component containing at least two elements of each of $S_{(0,0)}^{0,1}$ and $S_{(1,0)}^{1,1}$, let $\mathcal{C}_1$ be the analogous set for $S_{(0,0)}^{1,0}$ and $S_{(1,0)}^{1,1}$, $\mathcal{C}_2$ the analogous set for $S_{(0,0)}^{0,1}$ and $S_{(1,0)}^{0,0}$, and $\mathcal{C}_3$ the analogous set for $S_{(0,0)}^{1,0}$ and $S_{(1,0)}^{0,0}$. See Figure 2.3 for an illustration of the target sets.

Given a bond uv in the renormalised model, the relative states of $\mathcal{L}_u$ and $\mathcal{L}_v$ correspond (up to symmetry) to one of the $\mathcal{C}_i$. More specifically, if uv lies in one of the 2×2 squares tiling $\mathbb{Z}^2$ in the renormalised model, then the pair $(\mathcal{L}_u, \mathcal{L}_v)$ corresponds to the configurations $\mathcal{C}_0$, $\mathcal{C}_1$, $\mathcal{C}_2$, $\mathcal{C}_3$ with probabilities $(1 - \theta)^2$, $\theta(1 - \theta)$, $\theta(1 - \theta)$,

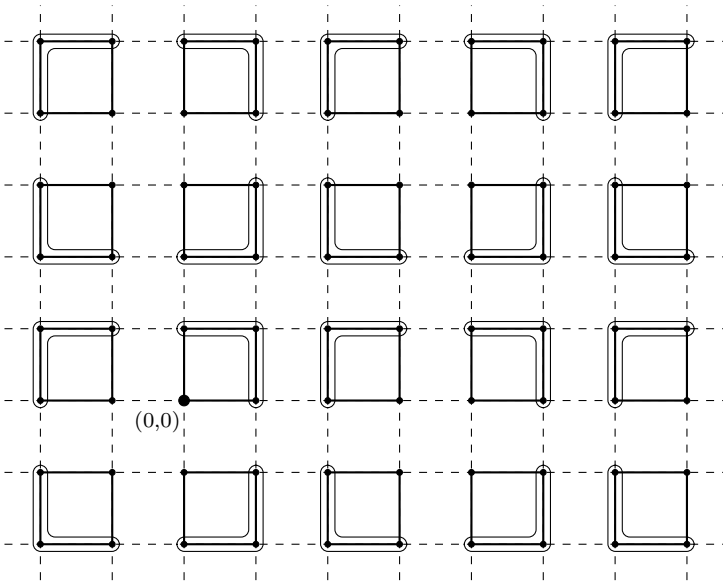


Figure 2.2: A portion of the graph $\mathbb{Z}^2$ with the regions corresponding to the $\mathcal{L}_u$ when $\theta = 0$ highlighted. Dashed edges represent edges open with probability p', other edges are open with probability p.

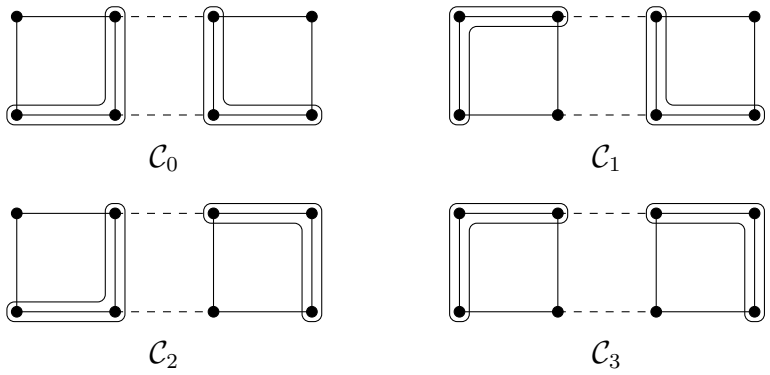


Figure 2.3: The configurations corresponding to the definitions of $\mathcal{C}_0$, $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$

and θ^2 respectively. Indeed, the default ($\theta = 0$) case is just $\mathcal{C}_0$ (up to symmetry) and then independently with probability θ we rotate each side by 180°. On the other hand, if uv joins two renormalised 2×2 squares, then in the default case we get a configuration that is $\mathcal{C}_3$ up to symmetry. Thus in general we obtain configurations $\mathcal{C}_0$, $\mathcal{C}_1$, $\mathcal{C}_2$, $\mathcal{C}_3$ with probabilities θ^2, $\theta(1-\theta)$, $\theta(1-\theta)$, and $(1-\theta)^2$ respectively.

Now, given probabilities p_0 and p_0', consider a 1-independent model on R in which edges inside $S_{(0,0)}$ and $S_{(1,0)}$ are open with probability at least p_0, and the two edges between $S_{(0,0)}$ and $S_{(1,0)}$ are open with probability at least p_0'. Let the random set of open edges be $E' \subseteq E$. For fixed $\theta \in [0,1]$, let p_1 and p_1' be such that, for all such

models on R,

$$p_1 \leqslant (1-\theta)^2\mathbb{P}(E' \in \mathcal{C}_0) + \theta(1-\theta)\big(\mathbb{P}(E' \in \mathcal{C}_1) + \mathbb{P}(E' \in \mathcal{C}_2)\big) + \theta^2\mathbb{P}(E' \in \mathcal{C}_3)$$

and

$$p_1' \leqslant \theta^2\mathbb{P}(E' \in \mathcal{C}_0) + \theta(1-\theta)\big(\mathbb{P}(E' \in \mathcal{C}_1) + \mathbb{P}(E' \in \mathcal{C}_2)\big) + (1-\theta)^2\mathbb{P}(E' \in \mathcal{C}_3).$$

Then by symmetry, if our original 1-independent model on $\mathbb{Z}^2$ lies in $\mathcal{D}_{\geqslant p_0, \geqslant p_0'}$ then the renormalised model lies in $\mathcal{D}_{\geqslant p_1, \geqslant p_1'}$.

We will iterate this, giving a $\theta \in [0,1]$ and a sequence of pairs (p_i, p_i'), $i = 0, \dots, k$ such that, for every model in $\mathcal{D}_{\geqslant p_i, \geqslant p_i'}$, the renormalised model lies in $\mathcal{D}_{\geqslant p_{i+1}, \geqslant p_{i+1}'}$. The theorem will follow if we can exhibit such a sequence with $p_0 = p_0' = 0.8457$ and $\min\{p_k, p_k'\} \geqslant 0.8639$, as there is then almost surely an infinite open component in the k-times renormalised model by the results of [17]. Note that at each stage we may assume that in the i times renormalised model each edge is open with probability *exactly* p_i or p_i' as appropriate. Indeed, independently deleting edges with the appropriate probabilities only makes the events $\mathcal{C}_i$ less likely.

After renormalising i times, we wish to minimise the right-hand sides of the two inequalities above over all models in $\mathcal{D}_{p_i, p_i'}$. As in the proof of Lemma 2.2, we relax the condition that we have a 1-independent model on R to obtain the following two linear programming problems.

$$\begin{aligned}
\text{Minimise: } & (1-\theta)^2 \textstyle\sum_{S \in \mathcal{C}_0} x_S + \theta(1-\theta)\big(\sum_{S \in \mathcal{C}_1} x_S + \sum_{S \in \mathcal{C}_2} x_S\big) + \theta^2 \sum_{S \in \mathcal{C}_3} x_S \\
\text{or: } & \theta^2 \textstyle\sum_{S \in \mathcal{C}_0} x_S + \theta(1-\theta)\big(\sum_{S \in \mathcal{C}_1} x_S + \sum_{S \in \mathcal{C}_2} x_S\big) + (1-\theta)^2 \sum_{S \in \mathcal{C}_3} x_S \\
\text{Subject to: } & x_S \geqslant 0, \\
& y_S = \textstyle\sum_{T \supseteq S} x_T, \\
& y_{S \cup \{e\}} = p_i \cdot y_S, \\
& y_{S \cup \{f\}} = p_i' \cdot y_S, \\
& y_\emptyset = 1,
\end{aligned}$$

where S runs over all subsets of E, e runs over all edges induced by $S_{(0,0)}$ or $S_{(1,0)}$ that are vertex-disjoint from all edges in S, and f runs over the edges $\{(1,0),(2,0)\}$ and $\{(1,1),(2,1)\}$ that are vertex-disjoint from all edges in S. The first of these optimisation problems yields a suitable value for p_{i+1} and the second yields a suitable value for p_{i+1}'.

Running the above LP problems with $\theta = 0.18$ using the `Gurobi` optimisation package we can find suitable values for p_i and p_i', which we checked by confirming that the dual LP problem was feasible. The results are listed in Table 2.1. For $k = 13$ we see that $p_k, p_k' \geqslant 0.8639$, and hence the model percolates by the results of [17]. □

i	p_i	p'_i
0	0.845700	0.845700
1	0.859167	0.829055
2	0.856981	0.831846
3	0.857370	0.831456
4	0.857391	0.831616
5	0.857546	0.831779
6	0.857826	0.832114
7	0.858365	0.832753
8	0.859396	0.833976
9	0.861358	0.836303
10	0.865058	0.840691
11	0.871911	0.848815
12	0.884171	0.863343
13	0.904695	0.887637
14	0.934851	0.923277

Table 2.1: Suitable values for p_i and p'_i (with $\theta = 0.18$)

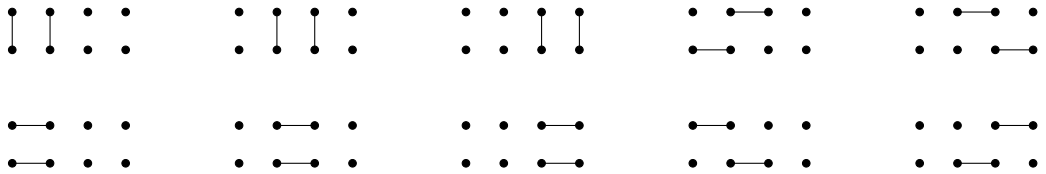
Of course, there is no need to use the same value of θ for every renormalisation in the proof of Theorem 1.3. Indeed, it seems rather unlikely that the optimal sequence of θ values is constant. We have, however, been unable to find a sequence that gives $p_{\max}(\mathbb{Z}^2) \leqslant 0.8456$. Note that taking $\theta = 0$ in the renormalisations, so that the $\mathcal{L}_u$ are as illustrated in Figure 2.2, yields only $p_{\max}(\mathbb{Z}^2) \leqslant 0.8463$. Moreover the rather simpler homogeneous renormalisation where we take $\mathcal{L}_u = S_u^{0,0}$ for all u, so that in the renormalised model we assume all edges are open with the same probability, gives $p_{\max}(\mathbb{Z}^2) \leqslant 0.8493$.

The renormalisation used in the proof of Theorem 1.3 can be combined with some ideas from [17] to give the following corollary.

Corollary 2.20. *In any model in $\mathcal{D}_{\geqslant 0.8457}(\mathbb{Z}^2)$ and for any $C > 0$, the probability that there is an open path crossing the rectangle $[0, Cn] \times [0, n]$ from left to right tends to 1 as $n \to \infty$. In particular, there are almost surely open cycles enclosing any bounded region of the plane.*

Proof. We recall from [17, Theorem 2] that a much simpler 2×2 renormalisation argument gives that if the original model is in $\mathcal{D}_{\geqslant 1-q}(\mathbb{Z}^2)$, then the renormalised model is in $\mathcal{D}_{\geqslant 1-10q^2}(\mathbb{Z}^2)$. Indeed, we just take the good event E_{uv} to be the event that there is a component meeting at least three out of the four vertices in each of

S_u and S_v. If this fails, we must have at least one of the following ten sets of edges closed.[2]



Hence, $\mathbb{P}(E_{uv}) \geqslant 1 - 10q^2$. Note that this lower bound on edge probabilities also applies to the renormalisation used in the proof of Theorem 1.3, since the good events there contain the good events here. Thus, once $p_k, p'_k > 0.9$ in the proof of Theorem 1.3 (see Table 2.1), we can inductively choose $p_{i+1} = p'_{i+1} = 1 - 10q_i^2$ where $q_k = 1 - \min\{p_k, p'_k\}$ and $q_i = 1 - p_i$ for all $i \geqslant k + 1$. It is clear that then $p_i \to 1$ rapidly as $i \to \infty$.

Now given any $\varepsilon > 0$ we can find some i_0 such that $p_i > 1 - \varepsilon / \lceil 2C + 1 \rceil$ for all $i \geqslant i_0$. Suppose $n \geqslant 2^{i_0}$ and pick i so that $2^i \leqslant n < 2^{i+1}$. Then in the i times renormalised model we have an open path from $(-1, 0)$ to $(\lceil 2C \rceil, 0)$ with probability at least $1 - \varepsilon$. It follows that in the original model there is an open path in $\mathbb{Z} \times \{0, \dots, n\}$ starting before $x = 0$ and ending at or after $x = \lceil 2C \rceil 2^i \geqslant Cn$ with probability at least $1 - \varepsilon$. Thus we have a path crossing $[0, Cn] \times [0, n]$ from left to right with probability at least $1 - \varepsilon$ as desired.

For the last part, fix a bounded set and let m_0 be large enough that $(-m_0, m_0)^2$ contains it. Then for all $m \geqslant m_0$, if there are open paths crossing the rectangles $[-2m, 2m] \times [m, 2m]$ and $[-2m, 2m] \times [-2m, -m]$ from left to right and the rectangles $[-2m, -m] \times [-2m, 2m]$ and $[m, 2m] \times [-2m, 2m]$ from top to bottom, then there is an open cycle enclosing the bounded set. By the above, in each case such a path is present with probability $1 - o(1)$ as $m \to \infty$, so such an open cycle exists with probability 1. It follows by a union bound that, almost surely, for every bounded set there is an open cycle enclosing it. □

As noted in [17], it follows from Corollary 2.20 that for every model in $\mathcal{D}_{\geqslant 0.8457}(\mathbb{Z}^2)$ the infinite component is almost surely unique: if not then pick two vertices in different infinite components, then there is an open cycle enclosing both of them and this connects the two infinite components. It also follows that there is no infinite component in the dual graph associated to such models.

[2] For each of the first three columns, having at least one of the two edges present in both upper and lower sets implies the corresponding 2×2 square has at least three vertices in the same component. The last two columns ensure that these components connect up.

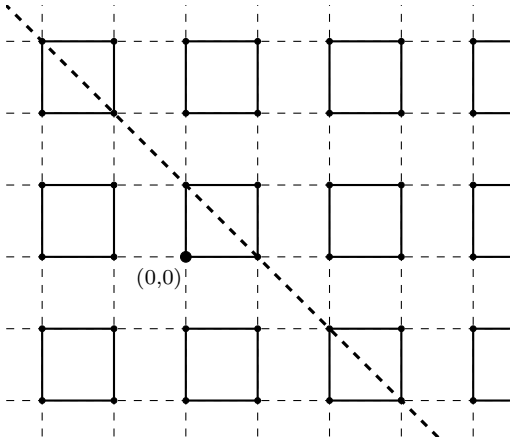


Figure 2.4: A portion of the graph $\mathbb{Z}^2$ and the line $x+y=1$. Dashed edges represent edges open with probability at least p', other edges are open with probability at least p.

In [17, Theorem 2] it was claimed that under every model in $\mathcal{D}_{\geqslant 0.8639}(\mathbb{Z}^2)$ the origin is in an infinite component with positive probability, but this result does not seem to follow from the proof given. However, we can now adapt the renormalisation used to prove Theorem 1.3 to show this claim holds for all models in $\mathcal{D}_{\geqslant 0.8459}(\mathbb{Z}^2)$.

Theorem 2.21. *In any model in $\mathcal{D}_{\geqslant 0.8459}(\mathbb{Z}^2)$ the origin is in an infinite open component with positive probability.*

Proof. We will begin by recursively defining a sequence of renormalisations. For $k \geqslant 0$, given the k-times renormalised model, first reflect the model in the line $x+y=1$ (i.e. map (x,y) to $(1-y,1-x)$) and then renormalise as in the proof of Theorem 1.3. The result is the $(k+1)$-times renormalised model. Note that the reflection preserves $\mathcal{D}_{\geqslant p,\geqslant p'}$ as the p-edges and p'-edges are mapped onto edges of the same type (see Figure 2.4).

We now make some observations about these renormalisations. Let $k \geqslant 0$ and let uv be an open edge in the k-times renormalised model. It is not difficult to see that in the restriction of the original model to the $2^k \times 2^{k+1}$ or $2^{k+1} \times 2^k$ rectangle corresponding to uv, there is a unique open component C^k_{uv} with the property that if the only open edges in this rectangle were those in C^k_{uv}, then the edge uv would still be open in the k-times renormalised model. We say that C^k_{uv} *gives rise to* uv. Suppose further that $k \geqslant 1$ and e is an open edge in the reflected $(k-1)$-times renormalised model that belongs to the open component in $S_u \cup S_v$ that intersects both S_u and S_v in large sets. It is also not difficult to see that the open component in the restriction of the

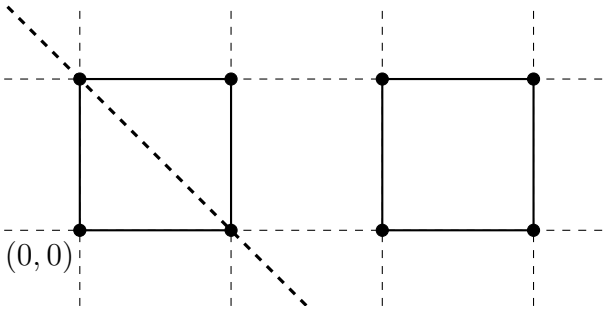


Figure 2.5: A portion of the reflected $(k-1)$-times renormalised model. If F_{k-1} holds, then the edge $(1,0)(1,1)$ is open, so if there is a component intersecting both $S_{(0,0)}$ and $S_{(1,0)}$ in large sets, then this component contains that edge.

original model to e that gives rise to e (i.e. $C^{k-1}_{e'}$ where e' is the reflection of e in the line $x + y = 1$) is contained in C^k_{uv}.

Next, we define events G_k, $k = 0, 1, 2, \ldots$ whose intersection we will show implies that the origin is in an infinite component in the original model. Let G_0 be the event that the edge from $(0,0)$ to $(1,0)$ is open in the original model, and for $k \geqslant 0$ let G_{k+1} be the event that, if the edge from $(1,0)$ to $(1,1)$ in the reflected k-times renormalised model were open, then $E_{(0,0)(1,0)}$ would hold. In other words, G_{k+1} is the event that adding the edge from $(1,0)$ to $(1,1)$ in the reflected k-times renormalised model gives a suitable connected component in $S_{(0,0)} \cup S_{(1,0)}$.

Let $F_k = \cap_{i=0}^k G_i$ for $k \geqslant 0$. It is straightforward to see by induction that F_k is contained in the event that the edge from $(0,0)$ to $(1,0)$ is open in the k-times renormalised model. We claim that, moreover, if F_k holds, then the origin of the original model is in $C^k_{(0,0)(1,0)}$.

We proceed by induction. The case $k = 0$ certainly holds. Hence, let $k \geqslant 1$ and assume that F_k holds. We will work in the reflected $(k-1)$-times renormalised model (see Figure 2.5). Since F_k is contained in F_{k-1}, the edge $(1,0)(1,1)$ is open in this model. Because we are working in the *reflected* k-times renormalised model the component in the restriction of the original model to $(1,0)(1,1)$ that gives rise to $(1,0)(1,1)$ is $C^{k-1}_{(0,0),(1,0)}$. By the induction hypothesis, this contains the origin in the original model. Since F_k holds, there is an open component in $S_{(0,0)} \cup S_{(1,0)}$ in the reflected $(k-1)$-times renormalised model that intersects both $S_{(0,0)}$ and $S_{(0,1)}$ in large sets. The edge $(1,0)(1,1)$ is open, so that component must include this edge. It follows from the observations above that $C^k_{(0,0)(1,0)}$ contains $C^{k-1}_{(0,0),(1,0)}$ and hence, in turn, the origin, as claimed.

As $C^k_{(0,0),(1,0)}$ strictly contains $C^{k-1}_{(0,0),(1,0)}$, we deduce that if $\cap_{k=0}^{\infty} G_k$ holds, then the origin is in an infinite component in the original model. Since

$$\mathbb{P}\Big(\bigcap_{k=0}^{\infty} G_k\Big) \geqslant 1 - \sum_{i=0}^{\infty} \mathbb{P}(G_k^c),$$

we seek to lower bound the probability of G_k for each k. Note that G_k depends only on the nine edges in $S_{(0,0)} \cup S_{(1,0)}$ other than that between $(1,0)$ to $(1,1)$, and is clearly an increasing event as a function of these edges. Hence, if we lower bound $\mathbb{P}(G_k)$ in $\mathcal{D}_{p,p'}$ we also have the same bound in $\mathcal{D}_{\geqslant p,\geqslant p'}$. It is straightforward to modify the linear programming problem from the proof of Theorem 1.3 to calculate such a lower bound for small k. For large k we note that clearly $\mathbb{P}(G_k^c) \leqslant q_k = 1 - p_k$, where the k-times renormalised model is in $\mathcal{D}_{\geqslant p_k,\geqslant p'_k}$, and we have the inequality $q_{k+1} \leqslant 10q_k^2$ as in the proof of Corollary 2.20.

To show that the origin is in an infinite component with positive probability, it is sufficient to show that $\sum_{k=0}^{\infty} \mathbb{P}(G_k^c) < 1$. Unfortunately, we cannot show this all the way down to $p = 0.8457$, but we show it does hold for $p = 0.8459$. Using the results from Table 2.2 we have

$$\begin{aligned}
\sum_{k=0}^{\infty} \mathbb{P}(G_k^c) &\leqslant \sum_{k=0}^{13} \mathbb{P}(G_k^c) + \sum_{k=14}^{\infty} q_k \\
&\leqslant 0.998359 + \sum_{i=1}^{\infty} 10^{2^i-1} q_{13}^{2^i} \\
&\leqslant 0.99836,
\end{aligned}$$

where we have used that $q_{13} \leqslant 0.0002$. Hence, $\mathbb{P}(\cap_{k=0}^{\infty} G_k) > 0$ and so the origin is in an infinite component with positive probability. □

2.6 Open problems

In this section we discuss some interesting related problems on 1-independent percolation models on hypercubes and lattices. We begin by restating a question first posed by Day, Falgas-Ravry, and Hancock [40].

Problem 2.22 ([40]). *For $n \geqslant 3$, what is the largest $p = p(n)$ for which there is a model in $\mathcal{D}_{\geqslant p}(Q_n)$ under which the hypercube is always disconnected?*

i	p_i	p'_i	$\mathbb{P}(G_k^c) \leqslant$
0	0.845900	0.845900	0.154100
1	0.859515	0.829480	0.096201
2	0.857661	0.832648	0.097540
3	0.858670	0.832999	0.096787
4	0.859879	0.834568	0.095945
5	0.862289	0.837404	0.094255
6	0.866795	0.842751	0.091100
7	0.875072	0.852561	0.085314
8	0.889637	0.869816	0.075168
9	0.913248	0.897752	0.058828
10	0.945814	0.936217	0.036503
11	0.978577	0.974824	0.014314
12	0.996611	0.996024	0.002247
13	0.999914	0.999899	0.000057

Table 2.2: Bounds used in Theorem 2.21 (with $\theta = 0.18$)

The simple construction used in the proof of Theorem 1.4 in [16] shows that this p satisfies $p \geqslant 1/2$, but is this best possible? The case $n = 2$ of Problem 2.22, where $Q_n = C_4$, was answered in [40] where they showed that the maximum is indeed $p = 1/2$ (this follows immediately from considering the expected number of edges in a model with edge probability greater than $1/2$).

It would also be interesting to determine the models above the threshold which minimise the probability that Q_n is connected. When $n = 2$, it is not difficult to show that these models are exactly those supported on subgraphs with either two or four edges, with

$$\mathbb{P}(\square) = 2p - 1, \quad \mathbb{P}(\,\vdots\,) = \mathbb{P}(\square) = q^2, \quad \mathbb{P}(\,\vdots\,) = \mathbb{P}(\square) = \alpha, \quad \mathbb{P}(\,\vdots\,) = \mathbb{P}(\square) = pq - \alpha,$$

where $q = 1 - p$ and $\alpha \in [0, pq]$.

The answer is already unknown for $n = 3$, although computer experiments suggest an answer for large p. Let the *signs model* on a graph G be the 1-independent percolation model where each vertex of G is independently assigned the sign '+' with probability θ and the sign '$-$' otherwise. An edge is open if both its endpoints have the same sign (so each edge is open with probability $p = \theta^2 + (1 - \theta)^2$). Computational evidence strongly suggests that this model minimises the probability that Q_3 is connected when $p > 0.55$. However, for smaller p there are models where connectivity is less likely, and indeed for $p < 0.516$ one can find models that are disconnected almost surely.

Conjecture 2.23. *For $p > 0.55$, among models in $\mathcal{D}_{\geqslant p}(Q_3)$ the signs model minimises the probability that Q_3 is connected.*

We also hypothesise that for general $n \geqslant 3$, the signs model minimises the probability that Q_n is connected when p is large enough. It should be noted that sadly none of the models listed above which minimise the probability of connectivity when $n = 2$ are signs models (unless $p = 1$).

Problem 2.24. *Is it true that for $n \geqslant 3$, if p is sufficiently close to 1, then among models in $\mathcal{D}_{\geqslant p}(Q_n)$ the signs model minimises the probability that Q_n is connected? If so, how close to 1 does p need to be?*

If the signs model is indeed optimal on Q_3, say for $p \geqslant 0.5720$, then this leads to improved bounds on p_{giant} and $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$ as follows. Start by applying the first step of the renormalisation in the proof of Theorem 2.3, with $k = 3$, to a model in $\mathcal{D}_{\geqslant p}(Q_n)$ to obtain a model on Q_{n-3} in which the probability that an edge is open given both its endpoints are open is at least

$$1 - \frac{2^{16}(1-p)^8}{\left(\left(1+\sqrt{2p-1}\right)^8 + \left(1-\sqrt{2p-1}\right)^8\right)^2}.$$

This expression is greater than p when $p \geqslant 0.5720$ so for such p we can apply this renormalisation repeatedly until we obtain a constant $I = I(p)$ and a model in $\mathcal{D}(Q_{n-I}, p_v, p_e)$ where $p_v > 0$ and $p_e > 18/19$ such that the existence of a giant component in this model implies the existence of a giant component in the original. It follows from Proposition 2.7 that there exists $C = C(p) > 0$ such that with high probability this model, and hence the original model, has a component containing at least a C proportion of its vertices. Thus, if the signs model is optimal on Q_3 for $p \geqslant 0.5720$, then $p_{\text{giant}} \leqslant 0.5720$. Similarly, $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n) \leqslant 0.5720$ under the same assumption.

We have shown that when each edge of the hypercube is open with probability at least 0.5847 (and n is large), there is a component containing a constant fraction of the vertices with high probability, so in particular $p_{\text{giant}} \leqslant 0.5847$. The best lower bound on this threshold is $1/2$ (using the model from the proof of Theorem 1.4 in [16] again). It was conjectured in [51] that this is tight, that is, $p_{\text{giant}} = 1/2$: they conjectured that for all $p > 1/2$, under every model in $\mathcal{D}_{\geqslant p}(Q_n)$ the hypercube contains a component of size at least $\left(\frac{1+\sqrt{2p-1}}{2} - o(1)\right) \cdot 2^n$ with probability $1 - o(1)$.

Problem 2.25. *What is the value of p_{giant}?*

It is not clear to us whether the threshold probabilities for always disconnecting Q_n (as in Problem 2.22) have a limit as $n \to \infty$, but if they do it would be interesting to know whether this limit differs from p_{giant}.

A very natural condition to impose on a percolation model on $\mathbb{Z}^n$ is that of translation invariance. Write $p_{\max}^T(\mathbb{Z}^n)$ for the translation invariant analogue of $p_{\max}(\mathbb{Z}^n)$, i.e.

$$\sup\big\{p : \text{some translation invariant model in } \mathcal{D}_{\geqslant p}(\mathbb{Z}^n) \text{ does not percolate}\big\}.$$

We will briefly discuss what is known in this setting for $n = 2$ and for general n. It would be interesting to see improvements on any of the bounds that follow.

Considering first the case $n = 2$, the bound

$$p_{\max}(\mathbb{Z}^2) \geqslant p_{\text{site}}^2 + \tfrac{1}{2}(1 - p_{\text{site}})$$

of Day, Falgas-Ravry, and Hancock [40] (where p_{site} denotes the critical probability for independent site percolation on $\mathbb{Z}^2$) arises from a translation invariant model. As discussed in Section 1.1, this therefore yields a rigorous bound of $p_{\max}^T(\mathbb{Z}^2) \geqslant 0.531136$ using the bound $p_{\text{site}} \geqslant 0.556$ of van den Berg and Ermakov [143] and a non-rigorous bound of $p_{\max}^T(\mathbb{Z}^2) \geqslant 0.554974$ using the widely believed estimate $p_{\text{site}} \approx 0.592746$ (see, for example, [157]). To our knowledge, this rigorous bound on $p_{\max}^T(\mathbb{Z}^2)$ is the best known. A better non-rigorous bound of $p_{\max}^T(\mathbb{Z}^2) \geqslant 0.5921$ follows from the high-confidence (p-value $< 10^{-11}$) bound on $p_{\max}(\mathbb{Z}^2)$ of Balister, Johnston, the author, and Scott [18]: the model from which that bound arises is also translation invariant.

Turning to upper bounds, an improved bound on $p_{\max}^T(\mathbb{Z}^2)$ can be obtained by modifying the linear programming problem from the proof of Theorem 1.3 to impose translation invariance, noting that after the first renormalisation the model will only be invariant under translations that preserve the structure of the p-edges and the p'-edges (i.e. translations by elements of $(2\mathbb{Z})^2$). Solving these LP problems with $\theta = 0.47$ and $p_0 = p_0' = 0.8265$ gives lower bounds on the renormalised edge probabilities as shown in Table 2.3. Since $p_{10}, p_{10}' > 0.8457$, it follows from Theorem 1.3 that the ten-times renormalised model percolates, and hence that $p_{\max}^T(\mathbb{Z}^2) \leqslant 0.8265$.

For general n (indeed, for all $n \geqslant 3$), to our knowledge the best known lower bound on $p_{\max}^T(\mathbb{Z}^n)$ is that arising from independent bond percolation. The threshold for independent bond percolation on $\mathbb{Z}^n$ is $(1 + o(1))/(2n - 1)$ (see, for example, [82]), so it is possible that $p_{\max}^T(\mathbb{Z}^n) = o(1)$. It would be very interesting to know whether this is the case. We have been unable to make use of translation invariance to show

i	p_i	p'_i
0	0.82650	0.82650
1	0.82831	0.82475
2	0.82833	0.82485
3	0.82844	0.82496
4	0.82865	0.82517
5	0.82906	0.82558
6	0.82984	0.82639
7	0.83137	0.82794
8	0.83430	0.83093
9	0.83986	0.83661
10	0.85018	0.84714

Table 2.3: Suitable values for p_i and p'_i (with $\theta = 0.47$) assuming the original model is translation invariant

an improved upper bound on $\lim_{n\to\infty} p^T_{\max}(\mathbb{Z}^n)$. In particular, adding translation invariance to the conditions of the LP problems from the proof of Lemma 2.2 does not yield a better bound, since the solutions found by `Gurobi` to these LP problems are already translation invariant.

Returning finally to the general (not necessarily translation invariant) setting, the rigorous bounds on $p_{\max}(\mathbb{Z}^2)$ and $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$ now stand at

$$\begin{aligned} 0.555197\ldots &\leqslant \quad p_{\max}(\mathbb{Z}^2) \quad &\leqslant 0.8457, \\ 0.535898\ldots &\leqslant \lim_{n\to\infty} p_{\max}(\mathbb{Z}^n) &\leqslant 0.5847. \end{aligned}$$

There remains a large gap between the upper and lower bounds in both cases and it would be interesting to reduce either of these gaps. Note that the non-rigorous bound of $p_{\max}(\mathbb{Z}^2) \geqslant 0.5921$ given in [18] shows that the lower bound for $p_{\max}(\mathbb{Z}^2)$ is unlikely to be correct, and suggests that the correct value of $p_{\max}(\mathbb{Z}^2)$ should be higher than $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$. We close by restating the fundamental Questions 1.1 and 1.2 from above, recalling that the authors of [51] conjecture that $p_{\max}(\mathbb{Z}^n) = 4 - 2\sqrt{3}$ for some $n \geqslant 3$.

Question 1.1 ([16])**.** *What is the value of* $p_{\max}(\mathbb{Z}^2)$*?*

Question 1.2 ([16])**.** *What is the value of* $\lim_{n\to\infty} p_{\max}(\mathbb{Z}^n)$*?*

Chapter 3

Game connectivity and adaptive dynamics

We start by recalling some standard definitions from the theory of games and introducing our notation. As usual, for $n \in \mathbb{N}$, we use $[n]$ as shorthand for the set $\{1, \dots, n\}$. For each $a \in \mathbb{N}^n$ and $i \in [n]$, we write a_{-i} for the element of $\mathbb{N}^{n-1}$ obtained by deleting the ith coordinate of a. In an abuse of notation, for $x \in \mathbb{N}$ and $a_{-i} \in \mathbb{N}^{n-1}$, we write (x, a_{-i}) for the element of $\mathbb{N}^n$ obtained by inserting x into the ith coordinate of a_{-i}.

A *game* is a tuple

$$\Big(n, \big(k_i\big)_{i\in[n]}, (\succsim_i)_{i\in[n]}\Big),$$

where $n \geqslant 2$ is an integer, $k_i \geqslant 2$ is an integer for each i, and for each i, $\succsim_i$ is a total preorder (i.e. a complete and transitive binary relation or, loosely, a total linear order in which ties are allowed) on $A \coloneqq \prod_{i\in[n]} [k_i]$. We say that $[n]$ is the *player set* of the game and that $[k_i]$ is the *action set* of player i. Elements of A are called *action profiles*, and $\succsim_i$ is known as i's *preference relation*. For each i, let $\succ_i$ denote the asymmetric part of $\succsim_i$.

We will discuss in detail how these games are played in Section 3.6 but, roughly, at each discrete time step every player simultaneously chooses an action, with the "aim" (in some sense) of player i being to maximise the resulting action profile $a \in A$ according to their preference relation $\succsim_i$. The information available to a player when they make this decision varies according to the specific dynamic under consideration. In this chapter a player will always know the current action profile and their own preference relation, but will not necessarily have access to historical action profiles or the preference relations of the other players.

Given an integer $n \geqslant 2$ and $\mathbf{k} = (k_1, \dots, k_n) \in \{2, 3, \dots\}^n$, we use $\mathcal{G}(n, \mathbf{k})$ to denote the set of all games with player set $[n]$ in which, for every $i \in [n]$, player i has action set $[k_i]$. For $i \in \mathbb{N}$ and $a_{-i} \in \prod_{j \in [n] \setminus \{i\}} [k_j]$, we write $L(a_{-i}) \coloneqq \{(x, a_{-i}) : x \in [k_i]\}$ for the set of action profiles in which all players other than player i play according to a_{-i}.

An action a_i of player i is a *better-response* than a_i' to a_{-i} if $(a_i, a_{-i}) \succ_i (a_i', a_{-i})$, that is, if player i would strictly prefer to play a_i than a_i' given that the other players play according to a_{-i}. An action a_i of player i is a *best-response* to a_{-i} if (a_i, a_{-i}) is $\succsim_i$-maximal in $L(a_{-i})$, that is, if $(a_i, a_{-i}) \succsim_i (x, a_{-i})$ for every $x \in [k_i]$. In words, a_i is a best-response to a_{-i} if player i has no better action than a_i given that the other players play according to a_{-i}. An action profile $a \in A$ is a *pure Nash equilibrium* if, for each player $i \in [n]$, a_i is a best-response to a_{-i}, that is, no player has an incentive to change their action given the actions of the other players.

A game is *generic* if for every i, a_{-i}, and distinct $(a_i, a_{-i}), (a_i', a_{-i}) \in L(a_{-i})$, either $(a_i, a_{-i}) \succ_i (a_i', a_{-i})$ or $(a_i', a_{-i}) \succ_i (a_i, a_{-i})$. In words, a game is generic if no player is ever indifferent between two of their possible actions given the actions of the other players.[1] We will focus almost entirely on generic games in this chapter.

The *better-response graph* of a game is the directed graph $(A, \to)$ whose vertex set is the set of action profiles A and whose directed edge set $\to$ is defined such that for $a, b \in A$,

> $a \to b$ if and only if there exists $i \in [n]$ such that $a_{-i} = b_{-i}$ and b_i is a better-response to a_{-i} than a_i.

In other words, the better-response graph is constructed by adding a directed edge from a to b if they differ in only one player's action, and that player strictly prefers b to a. The *best-response graph* of a game has vertex set A and directed edge set $\to$ such that for $a, b \in A$,

> $a \to b$ if and only if there exists $i \in [n]$ such that $a_{-i} = b_{-i}$ and b_i is both a better-response to a_{-i} than a_i, and a best-response to a_{-i}.

[1] Every game $(n, (k_i)_{i \in [n]}, (\succsim_i)_{i \in [n]})$ has a utility-based representation $(n, (k_i)_{i \in [n]}, (u_i)_{i \in [n]})$ with, for each player i, a utility function $u_i \colon A \to \mathbb{R}$ representing their preference relation $\succsim_i$. Genericity is equivalent to the condition that such a utility-based game has no payoff ties, i.e. that for each i and any distinct profiles a and a' that differ only in the ith coordinate, $u_i(a) \neq u_i(a')$. Note that any utility-based game can be made generic by perturbing the utility function.

Thus, the best-response graph is the subgraph of the better-response graph in which edges are kept when they represent not only a better-response, but a best-response.

Observe that a game's better- and best-response graphs are determined by the restrictions of $\succsim_i$ to $L(a_{-i})$ for each $i \in [n]$ and $a_{-i} \in \prod_{j \in [n] \setminus \{i\}} [k_j]$. Moreover, in the better-response graph of a *generic* game each $L(a_{-i})$ induces a transitive tournament, while in the best-response graph each $L(a_{-i})$ induces a star in which all edges are oriented towards the central vertex.

We now define various classes of games in terms of the connectivity properties of best- and better-response graphs.[2] As part of these definitions we will use standard terminology from the theory of directed graphs which we briefly recall here. Given a directed graph $(V, \to)$ on vertex set V with edge set $\to$, a vertex $v \in V$ is a *sink* if it has no outgoing edges, and a *non-sink* otherwise. Observe that pure Nash equilibria correspond to sinks in better- and best-response graphs. Similarly, a vertex $v \in V$ is a *source* if it has no incoming edges, and a *non-source* otherwise. For any pair of vertices $v, v' \in V$, we say that v can *reach* v' if there is a sequence $(v^1, \ldots, v^m)$ of vertices with $v^1 = v$ and $v^m = v'$ such that $v^i \to v^{i+1}$ for all $i \in [m-1]$; in this case we also say that the vertex v' can *be reached from* v. Note that every vertex can reach and be reached from itself. A *cycle* is a sequence $(v^1, \ldots, v^m)$ of distinct vertices that has length $m \geqslant 2$ and that satisfies $v^m \to v^1$ and $v^i \to v^{i+1}$ for all $i \in [m-1]$.

Definition 3.1. A game is *acyclic* if its best-response graph has no cycles. A game is *globally acyclic* if its better-response graph has no cycles.

Definition 3.2. A game is *weakly acyclic* if its best-response graph has the property that every vertex can reach a sink. A game is *globally weakly acyclic* if its better-response graph has the property that every vertex can reach a sink.

Observe that, by definition, a weakly acyclic game necessarily has at least one sink.

Definition 3.3. A game is *connected* if its best-response graph has at least one sink and the property that every non-sink can reach every sink. A game is *globally connected* if its better-response graph has at least one sink and the property that every non-sink can reach every sink.

Acyclicity, weak acyclicity, and their 'global' counterparts, have become standard concepts (see e.g. [50]) though they sometimes appear under different names in the

[2] See [34] for a 'flow'-based classification of games.

literature.[3] In this chapter, the terms acyclicity and weak acyclicity follow the terminology of Young [152], who introduced the concept of weak acyclicity to the literature on dynamics in games.

Acyclic games are a superset of the very widely studied class of potential games [107].[4] Potential games have been the subject of intense research, particularly because many dynamics are guaranteed to converge to a pure Nash equilibrium in such games (e.g. see [83, 125]). Weakly acyclic games are also very widely studied because (global) weak acyclicity is a necessary condition for the guaranteed convergence of (better-) best-response dynamics to a pure Nash equilibrium from any starting vertex (e.g. see [9, 50]).

The notion of connectedness, which is stronger than weak acyclicity, is one that we introduce in this work. Connected games have an important dynamical property not necessarily present in weakly acyclic games: because every non-sink can reach every sink along a best-response path it is, loosely, easy to move from each non-equilibrium point to *any* equilibrium point. Building on our work, Newton and Sawa [115] showed that this property can be used to greatly simplify the analysis of the long-term behaviours of various evolutionary dynamics, and characterised the pure Nash equilibria at which these dynamics spend most of their time in typical large games. Importantly, their arguments require genuine connectedness; weak acyclicity does not suffice. We discuss their work further in Section 3.6.

Remark 3.4. The following diagram summarises the relationships between the game classes above. Each inclusion is strict.

$$\begin{array}{ccccccc} \text{globally acyclic} & \Rightarrow & \text{acyclic} & \Rightarrow & \text{weakly acyclic} & \Rightarrow & \text{globally weakly acyclic} \\ & & & & \Uparrow & & \Uparrow \\ & & & & \text{connected} & \Longrightarrow & \text{globally connected} \end{array}$$

Chapter outline In Section 3.1 we state and discuss our main results on game connectivity. In Section 3.2 we state results on random subgraphs of directed Hamming graphs and explain how our game connectivity results follow. We also give an outline of the proof of our main theorem, and discuss how our results relate to previous work.

[3] In [142], for example, weak acyclicity is referred to as quasi-acyclicity.

[4] A (not necessarily generic) game $(n, (k_i)_{i\in[n]}, (\succsim_i)_{i\in[n]})$ is a *generalised ordinal potential game* if there exists a function $\rho\colon A \to \mathbb{R}$ such that for each $i \in [n]$ and each pair of distinct action profiles a and a' that differ in only the ith coordinate, $a \succ_i a'$ implies $\rho(a) > \rho(a')$. The game is an *ordinal potential game* or, simply, a *potential game* if, additionally, $\rho(a) > \rho(a')$ implies $a \succ_i a'$. Generalised ordinal potential games are precisely those that are globally acyclic [50, 107].

Section 3.3 is devoted to proving our main theorem, and Section 3.4 contains discussion and results on the tightness of various aspects of this theorem. In Section 3.5 we give the proof of our result on the prevalence of acyclic games. In Section 3.6 we describe and explain the consequences of our results for the study of dynamics in games, before closing in Section 3.7 with some open questions and directions for further study.

This chapter is based on joint work with Tom Johnston, Alex Scott, and Bassel Tarbush [87].

3.1 Game connectivity

The results we present in this section quantify the relative sizes of the game classes defined above for large games, where 'large' here means that the number of players is much bigger than the maximum number of actions available to each player. The following is our main result on game connectivity.

Theorem 3.5. *There exists $c > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2, 3, \dots\}^n$, if n is sufficiently large relative to $\max_i(k_i)$, then*

$$\frac{|\{g \in \mathcal{G}(n, \mathbf{k}) \colon g \text{ is generic and connected}\}|}{|\{g \in \mathcal{G}(n, \mathbf{k}) \colon g \text{ is generic and has a pure Nash equilibrium}\}|} \geqslant 1 - e^{-cn}.$$

Thus, connectedness is a ubiquitous property among large generic games that have a pure Nash equilibrium.[5]

Our condition for n to be 'sufficiently large' relative to $\max_i(k_i)$ is that all the k_i are at most $\delta\sqrt{n/\log(n)}$ for a suitable constant $\delta > 0$ (see Theorem 3.8 below). Theorem 3.5 has important implications for adaptive dynamics in games, on which we elaborate in Section 3.6.

[5] Rinott and Scarsini [119] show that

$$\lim_{n\to\infty} \frac{|\{g \in \mathcal{G}(n, \mathbf{k}) \colon g \text{ is generic and has a pure Nash equilibrium}\}|}{|\{g \in \mathcal{G}(n, \mathbf{k}) \colon g \text{ is generic}\}|} = 1 - e^{-1},$$

where this limit is uniform over all $\mathbf{k} = \mathbf{k}(n) \in \mathbb{N}^n$ with $k_i \geqslant 2$ for all i. Combined with Theorem 3.5, this implies that approximately 63% of all large generic games are connected.

In fact, building on work of Arratia, Goldstein, and Gordon [11], the authors of [119] show that the distribution of the number of pure Nash equilibria in games drawn uniformly at random from among all generic games is asymptotically Poisson(1) as the number of players with at least two actions gets large or as the number of actions gets large for at least two players. The distribution of pure Nash equilibria in random games had also previously been studied in [43, 63, 117, 138].

Our quantitative characterisation of the game classes defined at the start of this chapter is completed by the following result, which applies to all games, not just those that we consider 'large'.

Proposition 3.6. *There exists $c > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2, 3, \dots\}^n$, we have*

$$\frac{|\{g \in \mathcal{G}(n, \mathbf{k}) \colon g \text{ is generic and acyclic}\}|}{|\{g \in \mathcal{G}(n, \mathbf{k}) \colon g \text{ is generic and has a pure Nash equilibrium}\}|} \leqslant e^{-cn2^n}.$$

Together, Theorem 3.5 and Proposition 3.6 imply that there is a 'split' in large game properties: among large generic games that have a pure Nash equilibrium, acyclic and globally acyclic games are very rare, while connected games, weakly acyclic games, and their global counterparts, are very common. Note that since acyclic games are a superset of potential games, this also implies that potential games are very rare among large generic games that have a pure Nash equilibrium.

The results in this section will be proved as corollaries of results in the next section on the likelihood of analogous conditions holding in certain random subgraphs of directed Hamming graphs.

3.2 Directed grids

This section presents the main technical contributions of our work, including the statement of Theorem 3.8, a sketch of its proof, and a proof of its corollary, Theorem 3.5.

3.2.1 Game connectivity

Theorem 3.8, the main result of this section and the central technical contribution of this chapter, concerns the connectivity properties of random subgraphs of directed Hamming graphs. We first introduce our notation and then state the theorem, before explaining how our main result on game connectivity, Theorem 3.5, follows.

For $n \in \mathbb{N}$ and $\mathbf{k} = (k_1, \dots, k_n) \in \{2, 3, \dots\}^n$, the *Hamming graph* $H(n, \mathbf{k})$ is the graph with vertex set $V(n, \mathbf{k}) \coloneqq \prod_{i=1}^n [k_i]$ and edges between n-tuples precisely when they differ in exactly one coordinate. For $i \in [n]$, a *line of* $V(n, \mathbf{k})$ *in coordinate* i is a subset of $V(n, \mathbf{k})$ of size k_i whose elements pairwise differ in exactly the ith coordinate. A *line* of $V(n, \mathbf{k})$ is a subset which is a line of $V(n, \mathbf{k})$ in coordinate i for some i. Note that a line induces a complete subgraph of $H(n, \mathbf{k})$. The *directed*

Hamming graph $\vec{H}(n, \mathbf{k})$ is the simple directed graph formed by replacing each edge uv of $H(n, \mathbf{k})$ with directed edges $u \to v$ and $v \to u$.

Let $\vec{L}(n, \mathbf{k})$ be the random subgraph of the directed Hamming graph defined by independently and uniformly at random choosing a *winner* among the vertices of each line of $\vec{H}(n, \mathbf{k})$, and within that line keeping only those edges $u \to v$ whose endpoint, v, is the winner. Observe that in this random subgraph, each line induces a directed star in which all edges are oriented towards the winner. Our interest in $\vec{L}(n, \mathbf{k})$ stems from the following.

Remark 3.7. The graph $\vec{L}(n, \mathbf{k})$ has the same distribution as the best-response graph of a game drawn uniformly at random from amongst all generic games in $\mathcal{G}(n, \mathbf{k})$.

We will study these objects when n is large relative to $\max_i(k_i)$; our proof breaks down when this is not the case. We now state our main theorem.

Theorem 3.8. *For all $\varepsilon > 0$ there exist $c, \delta > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2, 3, \dots\}^n$, if $K \coloneqq \max_i(k_i)$ satisfies $K \leqslant \delta\sqrt{n/\log(n)}$, then with failure probability at most $\prod_{i=1}^n {k_i}^{-c}$, every vertex of $\vec{L}(n, \mathbf{k})$ can either be reached from at most $N \coloneqq (1+\varepsilon)K\log(K)$ vertices, or from every non-sink.*

The full proof of Theorem 3.8 will be postponed to Section 3.3 but we provide an outline of the proof in Section 3.2.3. In Section 3.4 we examine the tightness (or lack thereof) of various aspects of Theorem 3.8. In particular, we show that neither the failure probability nor the value of N can be significantly improved in general.

We remark that it is not particularly surprising that results such as Theorem 3.8 and, in turn, Theorem 3.5 should hold. Indeed, it is an easy exercise to show that a random n-vertex tournament is strongly connected (i.e. every vertex can reach every other vertex along a directed path) with exponentially small failure probability. Our base graph is, of course, rather sparser than the complete graph and our random subgraph model is different from independently orienting each edge of $\vec{H}(n, \mathbf{k})$. Nevertheless, directed connectivity properties of the type asserted by Theorem 3.8 should align with out intuition rather than being at odds with it.

We use the remainder of this subsection to explain how Theorem 3.5 follows from Theorem 3.8. To this end, we highlight the following corollary of Theorem 3.8, in which we denote by $R_{n,\mathbf{k}}$ the event that every non-sink in $\vec{L}(n, \mathbf{k})$ can reach every sink, and by $S_{n,\mathbf{k}}$ the event that $\vec{L}(n, \mathbf{k})$ has at least one sink.

Corollary 3.9. *There exist $c_0, c_1 > 0$ and $\delta \in (0,1]$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2,3,\dots\}^n$, if $K \coloneqq \max_i(k_i)$ is such that $K \leqslant \delta\sqrt{n/\log(n)}$, then*

(a) $\mathbb{P}(R_{n,\mathbf{k}}) \geqslant 1 - e^{-c_0 n}$,

(b) $\mathbb{P}(R_{n,\mathbf{k}} \,|\, S_{n,\mathbf{k}}) \geqslant 1 - e^{-c_1 n}$.

In the proof of Corollary 3.9 we will need that $\mathbb{P}(S_{n,\mathbf{k}})$ is bounded away from 0 for $n \geqslant 2$ and $\mathbf{k} \in \{2,3,\dots\}^n$. This follows from work in [119], but we include a short proof using the following standard result about factorial moments (see, for example, Corollary 1.11 in [21]). For a random variable X and an integer $r \geqslant 0$, the *rth factorial moment of X*, $\mathbb{E}_r[X]$, is defined as

$$\mathbb{E}_r[X] \coloneqq \mathbb{E}[X(X-1)\dots(X-r+1)].$$

Lemma 3.10. *Let $n \in \mathbb{N}$ and let X be a random variable that takes values in $\{0,1,\dots,n\}$. Then for each $k \in \{0,1,\dots,n\}$ we have*

$$\mathbb{P}(X = k) = \frac{1}{k!}\sum_{r=k}^{n}(-1)^{k+r}\frac{\mathbb{E}_r[X]}{(r-k)!}.$$

Moreover, the sum satisfies the alternating inequalities, that is, the sum upper bounds $\mathbb{P}(X = k)$ if truncated after an odd number of terms, and lower bounds it if truncated after an even number of terms.

We now prove Corollary 3.9.

Proof of Corollary 3.9. Let c and δ' be as given by Theorem 3.8 in the case $\varepsilon = 1$, and let δ be the minimum of 1 and δ'. Then for n, $\mathbf{k}$, and K as in the statement of the corollary, we have that with failure probability at most $\prod_{i=1}^n k_i^{-c}$ every vertex of $\vec{L}(n,\mathbf{k})$ can either be reached from at most $2K\log(K)$ vertices or from every non-sink. However, we have $2K\log(K) \leqslant K^2\log(K) \leqslant n$, so if this event holds then every non-sink can reach every sink, because all sinks can be reached from at least $n+1$ vertices. Finally, note that $\prod_{i=1}^n k_i^{-c} \leqslant 2^{-cn} \leqslant e^{-c_0 n}$ for some $c_0 > 0$, which proves part (a).

Next,

$$\mathbb{P}(R_{n,\mathbf{k}} \,|\, S_{n,\mathbf{k}}) = \frac{\mathbb{P}(R_{n,\mathbf{k}} \cap S_{n,\mathbf{k}})}{\mathbb{P}(S_{n,\mathbf{k}})} \geqslant \frac{\mathbb{P}(R_{n,\mathbf{k}}) - (1 - \mathbb{P}(S_{n,\mathbf{k}}))}{\mathbb{P}(S_{n,\mathbf{k}})} \geqslant 1 - \frac{e^{-c_0 n}}{\mathbb{P}(S_{n,\mathbf{k}})},$$

where we used part (a) in the final step. To prove part (b), it is enough to show that $\mathbb{P}(S_{n,\mathbf{k}}) \geqslant 1/2$.

Let $X_{n,\mathbf{k}}$ denote the number of sinks in $\vec{L}(n,\mathbf{k})$, then by Lemma 3.10 we have

$$1-\mathbb{P}(S_{n,\mathbf{k}})=\mathbb{P}(X_{n,\mathbf{k}}=0)\leqslant 1-\mathbb{E}[X_{n,\mathbf{k}}]+\tfrac{1}{2}\mathbb{E}_2[X_{n,\mathbf{k}}].$$

Clearly $\mathbb{E}[X_{n,\mathbf{k}}]=1$, so $\mathbb{P}(S_{n,\mathbf{k}})\geqslant 1-\frac{1}{2}\mathbb{E}_2[X_{n,\mathbf{k}}]$. We have

$$\mathbb{E}_2[X_{n,\mathbf{k}}]=\sum_{(v_1,v_2)}\mathbb{P}\big(v_1 \text{ and } v_2 \text{ are both sinks in } \vec{L}(n,\mathbf{k})\big),$$

where the sum is over all ordered pairs of distinct vertices in $V(n,\mathbf{k})$. If v_1 and v_2 share a line, then they cannot both be sinks, so the corresponding term in the sum is 0. If, on the other hand, they do not share a line, then they are sinks independently of one another, so the corresponding term in the sum is $\prod k_i^{-2}$. There are at most $\prod k_i^2$ terms in the sum, so $\mathbb{E}_2[X_{n,\mathbf{k}}]\leqslant 1$. It follows that $\mathbb{P}(S_{n,\mathbf{k}})\geqslant 1/2$. This completes the proof of part (b) of the corollary. □

As remarked above, $\vec{L}(n,\mathbf{k})$ has the same distribution as the best-response graph of a game drawn uniformly at random from among all generic games in $\mathcal{G}(n,\mathbf{k})$. Because our draws are uniform, we have that

$$\mathbb{P}(R_{n,\mathbf{k}}\,|\,S_{n,\mathbf{k}})=\frac{|\{g\in\mathcal{G}(n,\mathbf{k})\colon g \text{ is generic and connected}\}|}{|\{g\in\mathcal{G}(n,\mathbf{k})\colon g \text{ is generic and has a pure Nash equilibrium}\}|},$$

so Theorem 3.5 follows immediately from part (b) of Corollary 3.9.

3.2.2 Relation to previous work

In Section 3.4, we obtain the following as a corollary of Theorems 3.33 and 3.34. Here, for $k\geqslant 2$, we write $\mathbf{k}$ for the all k's vector of the appropriate length.

Corollary 3.11. *With failure probability $e^{-\Omega(n)}$, every non-sink in $\vec{L}(n,\mathbf{2})$ can reach every non-source. The same is true for $\vec{L}(n,\mathbf{3})$. Conversely, for each $k\geqslant 4$, the probability that there is a non-sink in $\vec{L}(n,\mathbf{k})$ which cannot reach every non-source tends to 1 as $n\to\infty$.*

Indeed, the $\mathbf{k}=\mathbf{2}$ part of the corollary follows immediately from Theorem 3.8 and the observation that every non-source in $\vec{L}(n,\mathbf{2})$ can be reached from at least two vertices. The $\mathbf{k}=\mathbf{3}$ part follows from a slightly adapted version of the proof of Theorem 3.8 (see Section 3.4). Note that the game property corresponding to the condition with which Corollary 3.11 is concerned is much stronger even than being connected.

Amiet, Collevecchio, Scarsini, and Zhong [7] derive connectivity properties of the graph $\vec{L}(n, \mathbf{2})$, though this is not the main focus of their paper. Their proofs show that for each fixed vertex x in $\vec{L}(n, \mathbf{2})$, x can reach every non-source in $\vec{L}(n, \mathbf{2})$ with failure probability $e^{-\Omega(n)}$ (see Theorem 4 in their paper).

Our results improve on this in two ways. Firstly, Theorem 3.8 and Corollary 3.11 assert that with high probability their respective reachability conditions hold for all vertices *simultaneously*, not just for a fixed vertex. For this reason, Corollary 3.11 is clearly stronger than the result of [7], even in the $\mathbf{k} = \mathbf{2}$ case. In particular, the result of [7] does not imply that almost all generic games with a pure Nash equilibrium are connected, or even weakly acyclic, in the $\mathbf{k} = \mathbf{2}$ case.

The second way our results improve on those in [7] is that clearly our results apply in cases where $k_i > 2$, and indeed k_i growing with n and varying amongst players in the case of Theorem 3.8. This is a non-trivial extension: the edges of $\vec{L}(n, \mathbf{2})$ are oriented independently of one another, and the proof in [7] relies heavily on this feature. Indeed, this independence allows them to deduce their result almost immediately from work of McDiarmid, Scott, and Withers [105] on the component structure of random subgraphs of the hypercube. In the more general case, edges in a line of $\vec{L}(n, \mathbf{k})$ are no longer independent of each other and for this reason our approach to the proof of Theorem 3.8 is different.

Finally, we remark that Corollary 3.11 shows that certain connectivity properties that hold with high probability in games with very few actions per player do not extend to games with more actions per player. The results of [7], which pertain to games with two actions per player, therefore cannot be used as a guide to what happens when there are more than two actions per player.

3.2.3 Outline of the proof of Theorem 3.8

We now provide a high-level explanation of the key elements of our proof of Theorem 3.8, the full version of which can be found in Section 3.3. At the heart of our proof is an argument showing the likely existence of a certain large strongly connected component in $\vec{L}(n, \mathbf{k})$; in other words, we find a large set of vertices which can all reach each other along directed paths. Indeed, we define a vertex of $\vec{L}(n, \mathbf{k})$ to be *good* if the number of lines that it wins is close to the expected number, and then show that with high probability every good vertex can reach all other good vertices along directed paths. Since most vertices win close to the expected number of lines,

this connects up a good proportion of the graph and it remains to 'plug in' the remaining, unusual, vertices. This part of our proof is based on work of McDiarmid, Scott, and Withers [105], who study the component structure of random subgraphs of the undirected hypercube graph.

Once we know that all good vertices are in the same strongly connected component, it is sufficient to show that every non-sink x can reach some good vertex u, and every vertex y which can be reached from more than $N = (1+\varepsilon)K\log(K)$ vertices can be reached from some good vertex v. Indeed, this yields a directed path from x to y via u and v, where we utilise the strongly connected component to get from u to v. The first step towards this is to 'establish a foothold' by showing that with high probability (a) every non-sink can in fact reach at least $n/2$ vertices, and (b) every vertex that can be reached from more than N vertices can in fact be reached from at least $n/2$ vertices. The inspiration for considering such an event comes from work of Bollobás, Gotsman, and Shamir [22]. The final step in the proof is then to show that if a vertex can reach or be reached from at least $n/2$ vertices, then it is very unlikely that none of these vertices is good.

3.2.4 Acyclicity of directed grids

Using the terminology and notation of Section 3.2.1 we can state the following strengthening of Proposition 3.6.

Proposition 3.12. *There exists $c > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2,3,\dots\}^n$, the probability that $\vec{L}(n,\mathbf{k})$ is acyclic is at most $\exp(-cnk^{n-2})$, where $k \coloneqq \min_i(k_i)$.*

The proof of Proposition 3.12 can be found in Section 3.5. Proposition 3.6 follows from Proposition 3.12 by the same reasoning with which we deduced Theorem 3.5 from Theorem 3.8: since $\vec{L}(n,\mathbf{k})$ has the same distribution as the best-response graph of a game drawn uniformly at random from all generic games in $\mathcal{G}(n,\mathbf{k})$, we have that $\mathbb{P}\big(\vec{L}(n,\mathbf{k}) \text{ is acyclic}\big)/\mathbb{P}(S_{n,\mathbf{k}})$ is equal to

$$\frac{|\{g \in \mathcal{G}(n,\mathbf{k})\colon g \text{ is generic and acyclic}\}|}{|\{g \in \mathcal{G}(n,\mathbf{k})\colon g \text{ is generic with a pure Nash equilibrium}\}|},$$

where $S_{n,\mathbf{k}}$ is the event that $\vec{L}(n,\mathbf{k})$ contains a sink, as in Section 3.2.1. Thus, since $k_i \geqslant 2$ for all i, Proposition 3.6 follows from Proposition 3.12 and the fact that $\mathbb{P}(S_{n,\mathbf{k}}) \geqslant 1/2$ for all n and $\mathbf{k}$ under consideration, as shown in the proof of Corollary 3.9 above.

3.3 Proof of Theorem 3.8

In this section we detail the proof of Theorem 3.8, which is restated below for convenience. We will continue to use the notation and terminology introduced in Section 3.2.

Theorem 3.8. *For all $\varepsilon > 0$ there exist $c, \delta > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2, 3, \dots\}^n$, if $K \coloneqq \max_i(k_i)$ satisfies $K \leqslant \delta\sqrt{n/\log(n)}$, then with failure probability at most $\prod_{i=1}^n k_i^{-c}$, every vertex of $\vec{L}(n, \mathbf{k})$ can either be reached from at most $N \coloneqq (1+\varepsilon)K\log(K)$ vertices, or from every non-sink.*

Throughout the proof of this theorem we will take $n \geqslant 2$ to be an integer, we will take $\mathbf{k} \in \{2, 3, \dots\}^n$, and we will let $K \coloneqq \max_i(k_i)$. Before going any further, it is worth remarking that by taking δ to be sufficiently small and imposing the condition $K \leqslant \delta\sqrt{n/\log(n)}$, we may assume that n is arbitrarily large in absolute terms (and indeed arbitrarily large relative to $K^2\log(K)$).

We will describe a probability $p_\varepsilon(n, \mathbf{k})$ with parameters $\varepsilon > 0$, n, and $\mathbf{k}$ as being *very small* if for all ε there exist $c_\varepsilon, \delta_\varepsilon > 0$ depending only on ε such that $p_\varepsilon(n, \mathbf{k}) \leqslant \prod_{i=1}^n k_i^{-c_\varepsilon}$ for all $K \leqslant \delta_\varepsilon\sqrt{n/\log(n)}$. For a probability $p(n, \mathbf{k})$ with no dependence on ε, the constants c_ε and δ_ε should be replaced by universal constants. Furthermore, we will say that $p(n, \mathbf{k})$ is *extremely small* if there exist $c, \delta_0 > 0$ such that for all $\delta \in (0, \delta_0)$, if $K \leqslant \delta\sqrt{n/\log(n)}$, then $p(n, \mathbf{k}) \leqslant e^{-cn\log(K)/\delta}$. Observe that every extremely small probability is also very small.

If $p_\varepsilon(n, \mathbf{k})$ or $p(n, \mathbf{k})$ is very or extremely small, we will say that the complementary probability is very or extremely high, respectively. We say that an event $F_\varepsilon(n, \mathbf{k})$ or $F(n, \mathbf{k})$ occurs *with very high probability* (wvhp) or *with extremely high probability* (wehp) if the probability that it occurs is very or extremely high respectively.

Given this terminology, Theorem 3.8 is equivalent to the statement that wvhp, every vertex of $\vec{L}(n, \mathbf{k})$ can either be reached from at most $(1+\varepsilon)K\log(K)$ vertices or from every non-sink. This motivates our definition of very small probabilities. Our definition of extremely small probabilities is motivated by the following lemma, which demonstrates that such probabilities are amenable to union bounds over $V(n, \mathbf{k})$.

Lemma 3.13. *If $p(n, \mathbf{k})$ is an extremely small probability, then for all fixed $a > 0$ the probability $K^{an} \cdot p(n, \mathbf{k})$ is also extremely small.*

Proof. Let c and δ_0 witness the fact that $p(n, \mathbf{k})$ is extremely small. Then for all $\delta \in (0, \delta_0)$, if $K \leqslant \delta\sqrt{n/\log(n)}$, then

$$K^{an} \cdot p(n, \mathbf{k}) \leqslant e^{an\log(K) - cn\log(K)/\delta} = e^{n\log(K)(a - c/\delta)}.$$

Let $\delta_0' \in (0, \delta_0)$ be small enough that $c/(2\delta_0') > a$, then for all $\delta \in (0, \delta_0')$ we have $a - c/\delta < -c/(2\delta)$, so letting $c' = c/2$ we see that c', δ_0' witness the fact that $K^{an} \cdot p(n, \mathbf{k})$ is extremely small. □

We will also make frequent use of the following simple result which follows from elementary analyses of the various cases.

Lemma 3.14. *The sum of two very small probabilities is very small and the sum of two extremely small probabilities is extremely small.*

Before starting the proof Theorem 3.8 in earnest, we record the following two standard results which will be useful at various points. For a discussion of these results (and much more), we refer the reader to [56]. First, in this chapter we will use the following version of the Chernoff bound.

Lemma 3.15 (Chernoff bound)**.** *Let $X_1, \ldots, X_n$ be independent Bernoulli random variables, let $X = \sum_{i=1}^n X_i$, and let $\mu = \mathbb{E}[X]$. Then for all $\varepsilon \geqslant 0$ we have*

$$\mathbb{P}(X \leqslant (1 - \varepsilon)\mu) \leqslant e^{-\varepsilon^2\mu/2}.$$

We will also need the following straightforward lemma, which can be proved via the first moment method.

Lemma 3.16. *Let $Z_1, \ldots, Z_m$ be non-negative integer valued random variables, and suppose that $\sum_{i=1}^m \mathbb{E}[Z_i] \leqslant p$ for some $p \in [0, 1]$. Then the probability that $Z_1 = \cdots = Z_m = 0$ is at least $1 - p$.*

3.3.1 Establishing a foothold

Throughout this section we let $\varepsilon > 0$ and set $N = (1+\varepsilon)K\log(K)$, as in the statement of Theorem 3.8. We will also assume (without loss of generality) that $k_1 \leqslant k_2 \leqslant \ldots \leqslant k_n$. This section is concerned with the following event.

Definition 3.17. Let A_ε be the event that the following two conditions are satisfied:

- all non-sinks can reach more than $n/2$ vertices; and

- every vertex that can be reached from more than N vertices can be reached from more than $n/2$ vertices.

We will "establish a foothold" by showing the following.

Lemma 3.18. *The event A_ε occurs with very high probability.*

Our proof of Lemma 3.18 follows an approach used by Bollobás, Gotsman, and Shamir [22] to study a random subgraph of $\vec{H}(n, \mathbf{2})$ with a similar distribution to that of $\vec{L}(n, \mathbf{2})$. Imitating those authors, for each $1 \leqslant m \leqslant \prod_{i=1}^{n} k_i$, define the random variable X_m to be the number of vertices of $\vec{L}(n, \mathbf{k})$ that can reach exactly m vertices (recall that every vertex can reach and be reached from itself). Analogously, let Y_m be the number of vertices that can be reached from exactly m vertices. Observe that A_ε can equivalently be defined as the event that $X_m = 0$ for all $2 \leqslant m \leqslant n/2$ and $Y_m = 0$ for all $N < m \leqslant n/2$.

Given a set $S \subseteq V(n, \mathbf{k})$, we say that a line of $V(n, \mathbf{k})$ in coordinate i is an *incomplete line of S* if its intersection with S has size other than 0 or k_i. If v can reach exactly m vertices, then running a depth first search from v gives a tree T with m vertices, in which all edges are oriented away from v and the winner of every incomplete line of T is in T. Similarly, if v can be reached from exactly m vertices, then we may build a tree T with m vertices where all the edges are oriented towards v and the winner of every incomplete line of T is outside of T. It follows that X_m and Y_m are bounded above by the number of pairs (v, T), where T is an m-vertex tree rooted at v of the appropriate one of these two forms.

We will use the following folklore result to upper bound the numbers of such trees in $H(n, \mathbf{k})$. A short combinatorial proof is given in [105].

Lemma 3.19. *If G is a graph with maximum degree Δ, then for each $m \in \mathbb{N}$ there are at most $(e\Delta)^{m-1}$ trees of order m in G that contain a given vertex.*

When applied to $H(n, \mathbf{k})$, Lemma 3.19 gives that there are at most $(enK)^{m-1}$ trees of order m in $H(n, \mathbf{k})$ that contain a given vertex. Lemma 3.18 follows from Lemma 3.14 and the next two lemmas, which handle the X_m and Y_m parts of the statement respectively.

Lemma 3.20. *With very high probability, $X_m = 0$ for all $2 \leqslant m \leqslant n/2$.*

Proof. We need to find universal $c, \delta > 0$ such that if $K \leqslant \delta\sqrt{n/\log(n)}$, then $X_m = 0$ for all $2 \leqslant m \leqslant n/2$ with probability at least $1 - \prod_{i=1}^n k_i^{-c}$. Thus, let $\delta > 0$ be small and assume that $K \leqslant \delta\sqrt{n/\log(n)}$.

Fix $2 \leqslant m \leqslant n/2$ and let T be a tree of order m in $H(n, \mathbf{k})$. Given the discussion preceding Lemma 3.19, we wish to upper bound the probability that the winner of every incomplete line of T is in T. To this end, it will be helpful to lower bound the number of lines containing exactly one vertex of T. Each vertex of T is in n lines, so there are mn pairs (u, ℓ) consisting of a vertex u in T and a line ℓ containing it. For each pair of distinct vertices u and v in T, if u and v are contained in some common line ℓ, then delete the pairs (u, ℓ) and (v, ℓ) from this set. Since any pair of vertices have at most one common line, this process removes at most $2\binom{m}{2}$ pairs from the set, and we deduce that there are at least $mn - m^2$ lines of $V(n, \mathbf{k})$ which contain exactly one vertex of T.

The winner of each of these lines is in T independently. Since we want to upper bound the probability that the winner of all of these lines is in T, we may assume that they are all in as low a coordinate direction as possible (recall that $k_1 \leqslant \ldots \leqslant k_n$ by assumption). At most m incomplete lines are in any given coordinate direction, so the probability that the winner of every incomplete line of T is in T is at most $\prod_{i=1}^{n-m} k_i^{-m}$.

By Lemma 3.19, the number of pairs (v, T) where $v \in V(n, \mathbf{k})$ and T is a tree of order m in $H(n, \mathbf{k})$ containing v is at most $(enK)^{m-1} \cdot \prod_{i=1}^n k_i$, so by the discussion before that lemma we have

$$\mathbb{E}[X_m] \leqslant \frac{(enK)^{m-1} \cdot \prod_{i=1}^n k_i}{\prod_{i=1}^{n-m} k_i^m} \leqslant \frac{K^m (enK)^{m-1} \cdot \prod_{i=1}^{n-m} k_i}{\prod_{i=1}^{n-m} k_i^m} \leqslant \frac{(enK^2)^m}{\prod_{i=1}^{n-m} k_i^{m-1}}.$$

Applying the fact that $m - 1 \geqslant m/2$ (since $m \geqslant 2$), we obtain

$$\mathbb{E}[X_m] \leqslant \frac{(enK^2)^m}{\prod_{i=1}^{n-m} k_i^{m/2}} \leqslant \left(\frac{enK^2}{\prod_{i=1}^{n-m} k_i^{1/2}}\right)^m \leqslant \prod_{i=1}^{n-m} k_i^{-m/3}$$

where the final inequality follows by taking δ small enough that $enK^2 \leqslant 2^{n/12}$, which is at most $\prod_{i=1}^{n-m} k_i^{1/6}$ since $m \leqslant n/2$.

Claim 3.21. *If δ is small enough, then $\prod_{i=1}^{n-m} k_i^{-m/3} \leqslant \prod_{i=1}^n k_i^{-1/2}$ for all $2 \leqslant m \leqslant n/2$.*

Proof. After rearranging, we need to show that $\prod_{i=n-m+1}^{n} k_i^{1/2} \leqslant \prod_{i=1}^{n-m} k_i^{m/3-1/2}$ for all $2 \leqslant m \leqslant n/2$. The left-hand side of this inequality is at most $K^{m/2}$ and the right-hand side is at least $2^{(n-m)(m/3-1/2)}$. Raising both sides to the power of $2/m$, it is sufficient that $K \leqslant 2^{(n-m)(2/3-1/m)}$. The right-hand side of this inequality is at least $2^{n/12}$, and we can take δ small enough that $K \leqslant 2^{n/12}$, so the claim is proved. □

Applying the claim, we have

$$\sum_{m=2}^{n/2} \mathbb{E}[X_m] \leqslant \frac{n}{2} \cdot \prod_{i=1}^{n} k_i^{-1/2}.$$

By taking δ to be sufficiently small we can ensure that this is at most $\prod_{i=1}^{n} k_i^{-c}$ for some $c > 0$. Lemma 3.16 now yields that $X_m = 0$ for all $2 \leqslant m \leqslant n/2$ with failure probability at most $\prod_{i=1}^{n} k_i^{-c}$, as required. □

The next lemma deals with the Y_m part of Lemma 3.18.

Lemma 3.22. *With very high probability, $Y_m = 0$ for all $N < m \leqslant n/2$.*

Proof. We need to show that there exist $c_\varepsilon, \delta_\varepsilon > 0$ depending only on ε such that if $K \leqslant \delta_\varepsilon \sqrt{n/\log(n)}$, then $Y_m = 0$ for all $N < m \leqslant n/2$ with failure probability at most $\prod_{i=1}^{n} k_i^{-c_\varepsilon}$. In fact, we will show a stronger failure probability of at most $e^{-c_\varepsilon n \log(K)}$. Thus, let $\delta_\varepsilon > 0$ be small and assume that $K \leqslant \delta_\varepsilon \sqrt{n/\log(n)}$.

We will employ a similar strategy to that used to prove Lemma 3.20. Fix $N < m \leqslant n/2$ and let T be a tree of order m in $H(n, \mathbf{k})$. We will upper bound the probability that the winner of every incomplete line of T is not in T using the lower bound of $mn - m^2$ on the number of incomplete lines of T from the proof of Lemma 3.20. The winner of each of these lines is in T independently, so the probability that all the winners are outside T is at most $(1 - 1/K)^{m(n-m)}$.

Hence, by Lemma 3.19 and the discussion preceding it, we have

$$\begin{aligned}\mathbb{E}[Y_m] &\leqslant K^n \cdot (enK)^{m-1} \cdot \left(1 - \frac{1}{K}\right)^{m(n-m)} \\ &\leqslant \left[enK^2 \cdot \left(K^{1/m}\left(1 - \frac{1}{K}\right)\right)^{n-m}\right]^m.\end{aligned}$$

Using that $m > N = (1+\varepsilon)K\log(K)$ and $1 + x \leqslant e^x$ for all x we have

$$K^{1/m}\left(1 - \frac{1}{K}\right) \leqslant K^{1/(1+\varepsilon)K\log(K)} e^{-1/K} = \exp\left(\frac{-\varepsilon}{(1+\varepsilon)K}\right). \tag{3.1}$$

Assuming that $\delta_\varepsilon \leqslant 1/2$, we have that $K^2 \leqslant n$. Applying this and $m \leqslant n/2$ yields

$$\mathbb{E}[Y_m] \leqslant \left[en^2 \cdot \exp\left(\frac{-\varepsilon(n-m)}{(1+\varepsilon)K}\right)\right]^m \leqslant \left[en^2 \cdot \exp\left(\frac{-\varepsilon n}{2(1+\varepsilon)K}\right)\right]^m.$$

By making δ_ε small enough that $\varepsilon n/(4(1+\varepsilon)K) \geqslant \log(en^2)$ and using the fact that $m \geqslant K\log(K)$, we have

$$\mathbb{E}[Y_m] \leqslant \exp\left(\frac{-\varepsilon mn}{4(1+\varepsilon)K}\right) \leqslant \exp\left(\frac{-\varepsilon n\log(K)}{4(1+\varepsilon)}\right).$$

Thus,

$$\sum_{N<m\leqslant n/2} \mathbb{E}[Y_m] \leqslant \frac{n}{2} \cdot \exp\left(\frac{-\varepsilon n\log(K)}{4(1+\varepsilon)}\right) \leqslant \exp\left(-\frac{\varepsilon}{8(1+\varepsilon)} n\log(K)\right)$$

for sufficiently large (depending only on ε) n. By making δ_ε sufficiently small relative to ε, we can ensure that we only need to consider values for n which are sufficiently large, and we find that $\sum_{N<m\leqslant n/2} \mathbb{E}[Y_m]$ is at most $e^{-c_\varepsilon n\log(K)}$ for $c_\varepsilon = \varepsilon/(8(1+\varepsilon))$. Lemma 3.16 now yields that $Y_m = 0$ for all $N < m \leqslant n/2$ with failure probability at most $e^{-c_\varepsilon n\log(K)}$, as required. □

With Lemmas 3.20 and 3.22 proved, Lemma 3.18 follows.

3.3.2 A strongly connected component

The next step in the proof of Theorem 3.8 will be to find a large strongly connected component in $\vec{L}(n,\mathbf{k})$. To this end, we make the following two definitions.

Definition 3.23. A vertex of $\vec{L}(n,\mathbf{k})$ is *good* if it wins at least $n/3K$ but at most $3n/4$ of its lines.

Definition 3.24. Let B be the event that every good vertex of $\vec{L}(n,\mathbf{k})$ can reach all other good vertices along directed paths, that is, all good vertices are in the same strongly connected component of $\vec{L}(n,\mathbf{k})$.

The main result of this section is that B occurs wehp.

Lemma 3.25. *The event B occurs with extremely high probability.*

Our proof of Lemma 3.25 is based on an approach taken by McDiarmid, Scott, and Withers [105] to analyse the component structure of the random subgraph of the n-dimensional hypercube in which edges are included independently with some constant probability. The proof will proceed via two auxiliary lemmas. The proof of

Lemma 3.26, and the proof of Lemma 3.25 given Lemmas 3.26 and 3.27, are straightforward generalisations of work in [105] and closely resemble arguments from Section 2.2.2, but we include them here nonetheless, for completeness. The proof of Lemma 3.27 bears some similarity to the corresponding argument from [105] in that it proceeds by considering many edge-disjoint paths between pairs of vertices. However, we need to introduce new ideas to overcome the fact that in that paper each edge appears with constant probability, which is not the case in our setting.

We now state our auxiliary lemmas.

Lemma 3.26. *With extremely high probability, every vertex in $V(n, \mathbf{k})$ has a neighbour in $H(n, \mathbf{k})$ which is good in $\vec{L}(n, \mathbf{k})$.*

Lemma 3.27. *With extremely high probability, every good vertex in $\vec{L}(n, \mathbf{k})$ can reach all good vertices within $H(n, \mathbf{k})$-distance 3 of it.*

Given Lemmas 3.26 and 3.27, Lemma 3.25 follows easily.

Proof of Lemma 3.25 given Lemmas 3.26 and 3.27. Let $\vec{F}$ be a realisation of $\vec{L}(n, \mathbf{k})$ for which the conclusions of Lemma 3.26 and Lemma 3.27 both hold, and note that this happens wehp by those two lemmas and Lemma 3.14. That is, let $\vec{F}$ be any outcome of $\vec{L}(n, \mathbf{k})$ in which every good vertex can reach all other good vertices at distance at most 3 from it in $H(n, \mathbf{k})$, and in which every vertex has a neighbour in $H(n, \mathbf{k})$ which is good. To prove Lemma 3.25, it suffices to show that for all good vertices x and y, there exists a directed path in $\vec{F}$ from x to y.

Let x and y be good vertices of $\vec{F}$ and choose a path $P = p_0 p_1 \dots p_t$ in $H(n, \mathbf{k})$ where $p_0 = x$ and $p_t = y$. If $t \leqslant 3$, then there exists a directed path from x to y in $\vec{F}$. Otherwise $p_2, p_3, \dots, p_{t-2}$ have good $H(n, \mathbf{k})$-neighbours $q_2, \dots, q_{t-2}$ respectively. Then x and q_2 are at distance at most 3, q_i and q_{i+1} are at distance at most 3 for all $2 \leqslant i \leqslant t-3$, and q_{t-2} and y are at distance at most 3, so $\vec{F}$ contains a path from x to y via $q_2, \dots, q_{t-2}$. □

Proof of Lemma 3.26. For a vertex $v \in V(n, \mathbf{k})$, let X be the number of lines that v wins in $\vec{L}(n, \mathbf{k})$ and note that v is good if and only if $n/3K \leqslant X \leqslant 3n/4$. Since X is a sum of n independent Bernoulli random variables and has mean $\mu = \sum_{i=1}^{n} 1/k_i \geqslant n/K$, by Lemma 3.15 we have

$$\mathbb{P}\Big(X < \frac{n}{3K}\Big) \leqslant \mathbb{P}\Big(X \leqslant \frac{\mu}{3}\Big) \leqslant e^{-2\mu/9} \leqslant e^{-2n/(9K)}.$$

Similarly, $n - X$ is a sum of n independent Bernoulli random variables and has mean $\mu' \geqslant n/2$, so by Lemma 3.15 again we have

$$\mathbb{P}\Big(X > 3n/4\Big) = \mathbb{P}\Big(n - X < n/4\Big) \leqslant \mathbb{P}\Big(n - X \leqslant \frac{\mu'}{2}\Big) \leqslant e^{-\mu'/8} \leqslant e^{-n/16}.$$

It follows (using $K \geqslant 2$) that v is good with failure probability at most $e^{-n/(10K)}$.

Now fix $u \in V(n, \mathbf{k})$ and pick one vertex other than u from each of the n lines containing it, say $v_1, \dots, v_n$. The v_i are distinct and no two of them share a line, so they are good independently of one another. Hence, the probability that u has no good $H(n, \mathbf{k})$-neighbour is at most $e^{-n^2/(10K)}$, so by a union bound over u, the probability that there exists a vertex with no good $H(n, \mathbf{k})$-neighbour is at most $K^n \cdot e^{-n^2/(10K)}$. For $0 < \delta \leqslant 1$, if $K \leqslant \delta\sqrt{n/\log(n)}$, then $K^2 \log(K) \leqslant \delta n$ and $n/K \geqslant K\log(K)/\delta \geqslant \log(K)/\delta$, so $e^{-n^2/(10K)}$ is extremely small. By Lemma 3.13, the same is true of $K^n \cdot e^{-n^2/(10K)}$, which completes the proof of the lemma. □

It remains to give the (slightly more involved) proof of Lemma 3.27.

Proof of Lemma 3.27. In this proof we will relabel $V(n, \mathbf{k})$ as $\prod_{i=1}^n \{0, \dots, k_i - 1\}$ in the natural way and will consider these vertices as elements of the vector space $\mathbb{R}^n$. We will write $\mathbf{e}_1, \dots, \mathbf{e}_n$ for the standard basis of this space.

We need to show that if $\delta > 0$ is sufficiently small, then whenever $K \leqslant \delta\sqrt{n/\log(n)}$, the probability that there exists a good vertex that cannot reach all other good vertices within $H(n, \mathbf{k})$-distance 3 of it is at most $e^{-cn\log(K)/\delta}$ for some universal $c > 0$. Thus, let $\delta > 0$ be small and assume that $K \leqslant \delta\sqrt{n/\log(n)}$. Note that since $n \geqslant 2$, by choosing δ small enough we may assume that $n/K^2 \geqslant \log(n)/\delta^2$ (and hence also n and n/K) is large in absolute terms.

Our proof will focus on pairs of vertices at distance exactly 3 from one another, and it will be clear how to adapt the argument to pairs at distance 1 or 2. Let u and v be vertices of $H(n, \mathbf{k})$ at distance 3 from each other. After relabelling, we may assume that $u = \mathbf{0}$ and $v = \mathbf{e}_1 + \mathbf{e}_2 + \mathbf{e}_3$. Fix subsets $A', B' \subseteq [n]$ of sizes at least $n/4$ and $n/3K$ respectively. Later, we will assume that u and v are good vertices, and will take A' to be the set of coordinate directions in which u *does not* win its line and B' to be the set of directions in which v *does* win its line. For each $i \in A'$ fix some $\alpha_i \in [k_i - 1]$; later we will take α_i to be the ith coordinate of the winner of the line through u in direction i. Now pick any $A \subseteq A' \setminus \{1, 2, 3\}$ and $B \subseteq B' \setminus \{1, 2, 3\}$

such that $|A| = \lceil n/5 \rceil$ and $|B| = \lceil n/4K \rceil$. Relabelling again, we may assume that $A, B \subseteq [\lfloor n/2 \rfloor]$.

Having fixed A, B, and the α_i, we will now define a certain type of path in $\vec{H}(n, \mathbf{k})$. First, let $i, j \in [n]$ be distinct with $i \in A$ and $j \in B$, then let $\beta \in [k_j - 1]$. Partition the interval of integers $[\lfloor n/2 \rfloor + 1, n]$ into disjoint intervals of integers I_1, I_2, and I_3 of as even lengths as possible (so that each interval has length approximately $n/3$). A directed path in $\vec{H}(n, \mathbf{k})$ from $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j$ to $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + v$ will be called an (i, j, β)-*path* if it has the following form: the path starts at $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j$ then follows a path of length 3 to $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \mathbf{e}_1$ in which the first and third edges are in a coordinate direction taken from I_1. That is, the path starts

$$\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j, \alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \gamma \mathbf{e}_\ell, \alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \gamma \mathbf{e}_\ell + \mathbf{e}_1, \alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \mathbf{e}_1, \tag{3.2}$$

for some $\ell \in I_1$ and $\gamma \in [k_\ell - 1]$. Next, the path follows a path of length 3 to $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \mathbf{e}_1 + \mathbf{e}_2$ in which the first and third edges are in a coordinate direction taken from I_2, before finally following a path of length 3 to $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \mathbf{e}_1 + \mathbf{e}_2 + \mathbf{e}_3 = \alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + v$ in which the first and third edges are in a coordinate direction taken from I_3.

Let $E^{(1)}_{(i,j,\beta)}$ be the event that there exists in $\vec{L}(n, \mathbf{k})$ a path of length 3 from $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j$ to $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \mathbf{e}_1$ in which the first and third edges are in a coordinate direction taken from I_1. That is, $E^{(1)}_{(i,j,\beta)}$ is the event that there is some $\ell \in I_1$ and some $\gamma \in [k_\ell - 1]$ such that all of the edges in the directed path (3.2) are present in $\vec{L}(n, \mathbf{k})$. Define $E^{(2)}_{(i,j,\beta)}$ and $E^{(3)}_{(i,j,\beta)}$ analogously for the second and third parts of the (i, j, β)-path. Note that there exists an (i, j, β)-path in $\vec{L}(n, \mathbf{k})$ if and only if all three of these events occur.

The event that $\vec{L}(n, \mathbf{k})$ contains a path of length 3 from $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j$ to $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \mathbf{e}_1$ in which the first and third edges are in a given coordinate direction ℓ has probability

$$\left(1 - \frac{1}{k_\ell}\right) \frac{1}{k_1} \frac{1}{k_\ell} \geqslant \frac{1}{2K^2}.$$

Indeed, the probability that $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j$ does not win its line in direction ℓ is $1 - 1/k_\ell$; the probability that $\alpha_i \mathbf{e}_i + \beta \mathbf{e}_j + \mathbf{e}_1$ wins its line in direction ℓ is $1/k_\ell$, the probability that the required edge in direction 1 is present is $1/k_1$, and these three events occur independently. Since the existence of such a path is independent for different ℓ, the failure probability of $E^{(1)}_{(i,j,\beta)}$ is at most

$$\prod_{\ell \in I_1} \left(1 - \frac{1}{2K^2}\right) \leqslant \left(1 - \frac{1}{2K^2}\right)^{n/7} \leqslant e^{-n/(14K^2)} < \frac{1}{2},$$

where we have used that $1 + x \leqslant e^x$ for all $x \in \mathbb{R}$ and that n/K^2 is large.

Similarly, $E^{(2)}_{(i,j,\beta)}$ and $E^{(3)}_{(i,j,\beta)}$ each occur with probability at least $1/2$. Moreover, it is not difficult to see that the sets of lines on whose presence each of these three events depend are pairwise disjoint, from which it follows that the events are independent. We deduce that $\vec{L}(n, \mathbf{k})$ contains an (i, j, β)-path with probability at least $1/8$.

Next, we will call a path in $\vec{H}(n, \mathbf{k})$ an *extended* (i, j, β)*-path* if it is an (i, j, β)-path extended by one vertex at the end to $\beta\mathbf{e}_j + v$. The line containing $\beta\mathbf{e}_j + v$ and $\alpha_i\mathbf{e}_i + \beta\mathbf{e}_j + v$ could never be used in any (i, j, β)-path, so the probability that $\vec{L}(n, \mathbf{k})$ contains an extended (i, j, β)-path is at least $1/(8k_i)$.

To conclude our definitions, a path in $\vec{H}(n, \mathbf{k})$ will be called an (i, j)*-path* if there exists $\beta \in [k_j - 1]$ for which it is an extended (i, j, β)-path extended by one vertex at the start to $\alpha_i\mathbf{e}_i$. Note that the line through $\alpha_i\mathbf{e}_i$ in coordinate j cannot be used in an extended (i, j, β)-path for any $\beta \in [k_j - 1]$. Hence, the probability that $\vec{L}(n, \mathbf{k})$ contains an (i, j)-path is the probability that the winner of the line through $\alpha_i\mathbf{e}_i$ in coordinate j is some $\alpha_i\mathbf{e}_i + \beta\mathbf{e}_j$ with $\beta \neq 0$, multiplied by the probability that $\vec{L}(n, \mathbf{k})$ contains an extended (i, j, β)-path for this β. By the above, this probability is at least

$$\frac{1}{8k_i}\left(1 - \frac{1}{k_j}\right) \geqslant \frac{1}{16k_i}.$$

Next, observe that every line that could possibly be used in an (i, j)-path identifies the set $\{i, j\}$. There are at least $|A|(|B| - 1)/2 \geqslant n^2/(50K)$ ways to choose $\{i, j\}$, so the probability that $\vec{L}(n, \mathbf{k})$ does not contain an (i, j)-path for any (i, j) is at most

$$\left(1 - \frac{1}{16k_i}\right)^{n^2/(50K)} \leqslant e^{-n^2/(800K^2)}$$

Observe also that the event that there exists an (i, j)-path in $\vec{L}(n, \mathbf{k})$ for some (i, j) is independent of the behaviour of any lines of $V(n, \mathbf{k})$ containing u or v.

This analysis holds for any choice of A' and B' containing at least $n/4$ and $n/3K$ directions respectively, so if u and v are good vertices, then we may take A' to be the set of coordinate directions in which u does not win its line, and B' to be the set of directions in which v does win its line. For each $i \in A'$, let $\alpha_i \in [k_i - 1]$ be the ith coordinate of the winner of the line through u in direction i. This means that the edges from u to $\alpha_i\mathbf{e}_i$ and from $\beta\mathbf{e}_j + v$ to v are both present for any choice of $i \in A$, $j \in B$, and $\beta \in [k_j - 1]$, so if there is no path from u to v in $\vec{L}(n, \mathbf{k})$, then

there is no (i,j)-path in $\vec{L}(n,\mathbf{k})$ for any (i,j). We have shown that this happens with probability at most $e^{-n^2/(800K^2)}$.

By a similar argument, the same holds for every pair of vertices at distance 1 or 2 from each other. A union bound yields that the probability that there exists a good vertex which cannot reach in $\vec{L}(n,\mathbf{k})$ some other good vertex within $H(n,\mathbf{k})$-distance 3 of it is at most $K^{2n}\cdot e^{-n^2/(800K^2)}$. Clearly $e^{-n^2/(800K^2)}$ is extremely small, so Lemma 3.13 implies that $K^{2n}\cdot e^{-n^2/(800K^2)}$ is also extremely small, which completes the proof of Lemma 3.27. □

With Lemma 3.27 proved, we have now completed the proof of Lemma 3.25. Note that at the end of the proof of Lemma 3.27 we used the full strength of the condition in Theorem 3.8 that $K \leqslant \delta\sqrt{n/\log(n)}$. Indeed, if $K = \omega(\sqrt{n/\log(n)})$, then $e^{-n^2/(800K^2)} = e^{-o(n\log(n))}$, so the union bound over all pairs of vertices at distance at most three from one another fails to produce a vanishing probability.

3.3.3 Accessing good vertices

In Section 3.3.2 we defined a vertex of $\vec{L}(n,\mathbf{k})$ to be *good* if it wins at least $n/3K$ but at most $3n/4$ of its lines. We will start by considering the following event.

Definition 3.28. Let C be the event that every vertex in $V(n,\mathbf{k})$ that can be reached from more than $n/2$ vertices in $\vec{L}(n,\mathbf{k})$ can be reached from a good vertex.

We show that C is very likely.

Lemma 3.29. *The event C occurs with extremely high probability.*

Proof. We need to show that if $\delta > 0$ is sufficiently small, then whenever $K \leqslant \delta\sqrt{n/\log(n)}$, we have $\mathbb{P}(\overline{C}) \leqslant e^{-cn\log(K)/\delta}$ for some $c > 0$. To this end, let $\delta > 0$ be small and assume that $K \leqslant \delta\sqrt{n/\log(n)}$. For each $y \in V(n,\mathbf{k})$, define C_y to be the event that either y can be reached from at most $n/2$ vertices in $\vec{L}(n,\mathbf{k})$, or y can be reached from a good vertex in $\vec{L}(n,\mathbf{k})$

Let $M = \lfloor n/3 \rfloor$. Fix y and let S be the set of trees of order M in $H(n,\mathbf{k})$ that contain y. For a tree $T \in S$, let C_T be the event that all the edges of T are oriented towards y in $\vec{L}(n,\mathbf{k})$ (so in particular, if C_T holds, then all vertices of T can reach y). If $\overline{C_y}$ occurs, then y can be reached from more than $n/2 \geqslant M$ vertices, so C_T occurs for

some $T \in S$. Hence, by a union bound

$$\mathbb{P}\big(\overline{C_y}\big) = \mathbb{P}\left(\overline{C_y} \cap \left(\bigcup_{T \in S} C_T\right)\right) \leqslant \sum_{T \in S} \mathbb{P}\big(\overline{C_y} \cap C_T\big). \tag{3.3}$$

For fixed $T \in S$, if $\overline{C_y}$ and C_T both occur, then no vertex in T is good. Since T contains exactly M vertices, each of its vertices can be assigned a set of $n - M = \lceil 2n/3 \rceil$ lines which contain that vertex and no other vertex of T (so that the numbers of lines that each vertex of T wins out of the $\lceil 2n/3 \rceil$ assigned to it are independent). For each vertex v of T, let X_v be the number of the $\lceil 2n/3 \rceil$ lines assigned to v in which v is the winner, so that X_v is a sum of $\lceil 2n/3 \rceil$ independent Bernoulli random variables and has mean $\mu \geqslant 2n/3K$.

By Lemma 3.15,

$$\mathbb{P}\left(X_v \leqslant \frac{n}{3K}\right) \leqslant \mathbb{P}\left(X_v \leqslant \frac{\mu}{2}\right) \leqslant e^{-\mu/8} \leqslant e^{-n/(12K)}.$$

Meanwhile, the number of the lines assigned to v in which v is not the winner, $\lceil 2n/3 \rceil - X_v$, is a sum of $\lceil 2n/3 \rceil$ independent Bernoulli random variables with mean $\mu' \geqslant n/3$. Hence, by Lemma 3.15 again,

$$\mathbb{P}\Big(X_v > \frac{5n}{12}\Big) \leqslant \mathbb{P}\Big(\Big\lceil \frac{2n}{3} \Big\rceil - X_v \leqslant \frac{n}{4}\Big) \leqslant \mathbb{P}\Big(\Big\lceil \frac{2n}{3} \Big\rceil - X_v \leqslant \frac{3\mu'}{4}\Big) \leqslant e^{-\mu'/32} \leqslant e^{-n/96}.$$

If $n/3K < X_v \leqslant 5n/12$, then in total v wins at least $n/3K$ and at most $5n/12 + \lfloor n/3 \rfloor \leqslant 3n/4$ of its lines, and hence is a good vertex. By the above, for each vertex v of T, we have $n/3K < X_v \leqslant 5n/12$ with failure probability at most $e^{-n/(50K)}$ (for $n \geqslant 101$). These events are independent for different v, so at least one vertex of T is good with failure probability at most $e^{-Mn/(50K)} \leqslant e^{-n^2/(200K)}$ (by choosing δ small enough, we can ensure $M \geqslant n/4$). This is therefore an upper bound on $\mathbb{P}(\overline{C_y} \cap C_T)$. It follows from Lemma 3.19 that $|S| \leqslant (enK)^{M-1}$, so by (3.3) we have

$$\begin{aligned}
\mathbb{P}\big(\overline{C_y}\big) &\leqslant (enK)^{M-1} \cdot e^{-n^2/(200K)} \\
&\leqslant \exp\left(\frac{n \log(enK)}{3} - \frac{n^2}{200K}\right) \\
&= \exp\left(\frac{n^2}{K}\left(\frac{K \log(enK)}{3n} - \frac{1}{200}\right)\right).
\end{aligned}$$

If δ is small enough, then n is large relative to $K \log(enK)$, so this probability is at most $e^{-n^2/(300K)}$.

Finally, since $C = \bigcap_y C_y$, by a union bound we have

$$\mathbb{P}(\overline{C}) \leqslant K^n \cdot e^{-n^2/(300K)}.$$

Clearly $e^{-n^2/(300K)}$ is extremely small, so by Lemma 3.13 the same is true of $K^n \cdot e^{-n^2/(300K)}$, and the lemma follows. □

Reversing orientations in the argument above shows that wehp every vertex in $\vec{L}(n, \mathbf{k})$ that can *reach* more than $n/2$ vertices can *reach* a good vertex. We record the analogues of Definition 3.28 and Lemma 3.29 in this setting.

Definition 3.30. Let D be the event that every vertex in $\vec{L}(n, \mathbf{k})$ that can reach more than $n/2$ vertices can reach a good vertex.

Lemma 3.31. *The event D occurs with extremely high probability.*

We are now ready to put everything together to prove the main theorem.

Proof of Theorem 3.8. Let $n \geqslant 2$ be an integer, let $\mathbf{k} \in \{2, 3, \dots\}^n$, and let $\varepsilon > 0$. Define $K = \max_i(k_i)$ and $N = (1+\varepsilon)K\log(K)$. Let E_ε be the event that every vertex of $\vec{L}(n, \mathbf{k})$ can either be reached from at most N vertices or can be reached from every non-sink; we want to show that E_ε occurs wvhp.

Let events A_ε, B, C, and D be as above, and suppose that they all occur simultaneously in $\vec{L}(n, \mathbf{k})$. Let $x, y \in V(n, \mathbf{k})$ where x is a non-sink and y can be reached from more than N vertices. Since A_ε occurs, x and y can reach and be reached from more than $n/2$ vertices respectively. Thus, since C occurs, there exists a good vertex $v \in V(n, \mathbf{k})$ which can reach y in $\vec{L}(n, \mathbf{k})$. Next, since D occurs, there exists a good vertex $u \in V(n, \mathbf{k})$ which can be reached from x in $\vec{L}(n, \mathbf{k})$. Since B occurs, all good vertices of $\vec{L}(n, \mathbf{k})$ are in the same strongly connected component, so in particular u can reach v. It follows that there is a directed walk from x to y in $\vec{L}(n, \mathbf{k})$ via u and v, that is, x can reach y. In other words, if A_ε, B, C, and D occur, then so does E_ε. By Lemmas 3.25, 3.29, and 3.31, each of B, C, and D occurs wehp, so in particular wvhp, and A_ε occurs wvhp by Lemma 3.18. Hence, by (repeated applications of) Lemma 3.14, we conclude that E_ε occurs wvhp, as required. □

3.4 The tightness of Theorems 3.5 and 3.8

In this section we discuss to what extent various aspects of Theorem 3.5 and Theorem 3.8 are tight. First, with regards to the relationship between n and K, it is

entirely possible that this condition could be weakened considerably while still allowing results in the spirit of Theorem 3.5 and Theorem 3.8. However, the condition given in Theorem 3.8 seems to be at the limit of our methods, and a new approach would be needed to improve it. See also point (i) in Section 3.7.

Next, with regards to the tightness of the failure probability in the theorems, first note that the failure probability in Theorem 3.8 carries through to Theorem 3.5 by the same argument as in Section 3.2.1. The following theorem shows that this slightly stronger lower bound is tight up to the value of the exponent c. In fact, the theorem shows that even if Theorem 3.5 were weakened to only consider weakly acyclic games rather than connected games, the probability inherited from Theorem 3.8 would still be tight up to the value of the constant in the exponent.

Theorem 3.32. *There is a constant $c' > 0$ such that*

$$\frac{|\{g \in \mathcal{G}(n,\mathbf{k})\colon g \text{ is generic and weakly acyclic}\}|}{|\{g \in \mathcal{G}(n,\mathbf{k})\colon g \text{ is generic and has a pure Nash equilibrium}\}|} \leqslant 1 - \prod_{i=1}^{n} k_i^{-c'}$$

for all integers $n \geqslant 4$ and all $\mathbf{k} \in \{2,3,\dots\}^n$.

Proof. Let n and $\mathbf{k}$ be as in the statement of the theorem. Define a 4-cycle in a subgraph of $\vec{H}(n,\mathbf{k})$ to be *sticky* if each of its vertices can only reach the other vertices of the 4-cycle. The probability that a given sticky 4-cycle (where the edges in the cycle are the first and second coordinate directions, say) appears in $\vec{L}(n,\mathbf{k})$ is $k_1^{-2}k_2^{-2}\prod_{i=3}^{n} k_i^{-4} \geqslant \prod_{i=1}^{n} k_i^{-4}$. Since $n \geqslant 4$, there is a vertex none of whose lines intersect this sticky 4-cycle. The event that this vertex is a sink, which occurs with probability $\prod_{i=1}^{n} k_i^{-1}$, is therefore independent of whether or not the sticky 4-cycle appears, so with probability at least $\prod_{i=1}^{n} k_i^{-5}$ there is both a sink and a sticky 4-cycle in $\vec{L}(n,\mathbf{k})$.

The vertices of a sticky 4-cycle cannot reach a sink, so the result now follows from arguments similar to those used to prove Corollary 3.9. □

One might ask whether taking a larger value for N in the statement of Theorem 3.8 would allow a significantly smaller failure probability, but a similar argument to the proof of Theorem 3.32 shows that it would not. Indeed, there is some $c' > 0$ such that if n is large relative to $\max_i(k_i)$, then with probability at least $1 - \prod_{i=1}^{n} k_i^{-c'}$ there is a vertex in $\vec{L}(n,\mathbf{k})$ which can be reached from $\prod_{i=1}^{n-1} k_i$ vertices but not from every non-sink.

This follows from an argument similar to the above: suppose that the desired sticky 4-cycle has vertices $(1,1,1,\dots,1)$, $(1,2,1,\dots,1)$, $(2,1,1,\dots,1)$, and $(2,2,1,\dots,1)$, then the subgraph G of $\vec{L}(n,\mathbf{k})$ induced on $\prod_{i=1}^{n-1}[k_i]\times\{2\}$ has the same distribution as $\vec{L}(n-1,(k_1,\dots,k_{n-1}))$, and behaves independently of whether the desired sticky 4-cycle appears or not. Applying Theorem 3.8 to G and using work of [119], one can show that there exists a universal $p>0$ such that if n is large enough relative to $\max_i(k_i)$, then with probability at least p, G contains exactly one sink and this can be reached from every vertex in G. It follows that there exists $c'>0$ such that if n is large relative to $\max_i(k_i)$, then with probability at least $\prod_{i=1}^{n}k_i^{-c'}$ there is a vertex in $\vec{L}(n,\mathbf{k})$ which can be reached from $\prod_{i=1}^{n-1}k_i$ vertices but not from every non-sink.

Finally, to what extent is it possible to take a smaller N in the statement of Theorem 3.8? It turns out that, for large K, the value of N cannot be substantially improved as a function of K: it is possible to take r not much smaller than $\log(K-1)$ in the following theorem,[6] so one cannot hope for a value of N any better than $K\log(K)-O(\log(K))$. Here we let $\mathbf{K}=(K,\dots,K)$ denote the all K's vector of the appropriate length.

Theorem 3.33. *There is a constant $c>0$ such that for all integers $n\geqslant 2$, $2\leqslant K\leqslant\sqrt{n}$, and*

$$1\leqslant r\leqslant\frac{\log(K-1)}{(K-1)(\log(K)-\log(K-1))},$$

the probability that there is a vertex in $\vec{L}(n,\mathbf{K})$ which can be reached from exactly $r(K-1)+1$ vertices is at least $1-c/n$.

Proof. Let $f(K)$ denote the expression upper bounding r in the theorem. One can show that this is increasing for $K\geqslant 2$ and that $f(\sqrt{n})\leqslant n$ for $n\geqslant 2$, so letting n, K, and r be as in the statement, we have $r\leqslant n$. Let X_a be the indicator random variable of the event that $a\in V(n,\mathbf{K})$ wins exactly r of its lines and every vertex on those r lines except a is a source, and write $X=\sum_{a\in[K]^n}X_a$. We wish to upper bound the probability that $X=0$, for which we will use a second moment calculation.

First, note that X_a and X_b are independent if the Hamming distance between a and b (i.e. the number of coordinates on which a and b differ), denoted by $d(a,b)$, is at least

[6] By applying the mean value theorem to log one can show that $\log(K)-\log(K-1)=1/K+O(1/K^2)$.

four. It follows that

$$\begin{aligned}\mathbb{E}[X^2] &= \sum_{a\in[K]^n}\sum_{b\in[K]^n}\mathbb{P}(X_aX_b=1)\\ &\leqslant \sum_{a\in[K]^n}\sum_{b\in[K]^n}\mathbb{P}(X_a=1)\mathbb{P}(X_b=1)+\sum_{a\in[K]^n}\sum_{b\in[K]^n:d(a,b)\leqslant 3}\mathbb{P}(X_a=1)\\ &\leqslant \mathbb{E}[X]^2+K^3n^3\,\mathbb{E}[X]\\ &\leqslant \mathbb{E}[X]^2+n^{9/2}\,\mathbb{E}[X].\end{aligned}$$

Hence, to apply Chebyshev's inequality, we need to show that $\mathbb{E}[X]$ grows more quickly than $n^{9/2}$. We have

$$\begin{aligned}\mathbb{E}[X] &= K^n\binom{n}{r}\frac{1}{K^r}\left(1-\frac{1}{K}\right)^{n-r+r(n-1)(K-1)}\\ &\geqslant \left(\frac{n}{Kr}\right)^r\left[K\left(1-\frac{1}{K}\right)^{1+r(K-1)}\right]^n\\ &\geqslant \left(\frac{\sqrt{n}}{r}\right)^r\left[K\left(1-\frac{1}{K}\right)^{1+r(K-1)}\right]^n,\end{aligned}$$

where we have used $\binom{n}{r}\geqslant(n/r)^r$ in the second line and $K\leqslant\sqrt{n}$ in the last line.

We will analyse the two terms in this product separately. For fixed n, the first term, $(\sqrt{n}/r)^r$, is increasing for $r\in[0,\sqrt{n}/e]$. Since $f(\sqrt{n})\leqslant\sqrt{n}/e$ for all $n\geqslant 2$, it follows that this term is always at least $\sqrt{n}\geqslant 1$, and if $r\geqslant 11$, then it is at least $n^{11/2}/11^{11}$.

Turning to the second term, it is straightforward to check that if $f(K)=r$, then the expression in square brackets is equal to 1, and that for fixed $K\geqslant 2$ this expression is strictly decreasing for $r\in[1,f(K)]$. It follows that the second term is always at least 1. There are no integer solutions K to $f(K)=r$ for any $r\in[10]$, so for $r\leqslant 10$ the expression in square brackets is always strictly greater than 1. Moreover, for fixed r we have $K(1-1/K)^{1+r(K-1)}\to\infty$ as $K\to\infty$, so in fact there exists some universal $\varepsilon>0$ such that the second term is at least $(1+\varepsilon)^n$ whenever $r\leqslant 10$.

Combining, we have $\mathbb{E}[X]\geqslant\min\left\{(1+\varepsilon)^n, n^{11/2}/11^{11}\right\}$ for all admissible n, K, and r. By Chebyshev's inequality, this yields

$$\mathbb{P}\left(|X-\mathbb{E}[X]|\geqslant\tfrac{1}{2}\mathbb{E}[X]\right)\leqslant 4\cdot\frac{\mathbb{E}[X^2]-\mathbb{E}[X]^2}{\mathbb{E}[X]^2}\leqslant\frac{4n^{9/2}}{\min\left\{(1+\varepsilon)^n, n^{11/2}/11^{11}\right\}},$$

and since $\mathbb{E}[X]>0$ this is, in turn, an upper bound on $\mathbb{P}(X=0)$. Clearly there is a constant $c'>0$ such that this upper bound is at most c'/n for all sufficiently large n,

and we can choose $c > 0$ large enough to accommodate the (finitely many) remaining cases. □

If, however, we are prepared to allow the exponent in the failure probability to depend on K, then we can adapt the proof of Theorem 3.8 to slightly improve the value of N.

Theorem 3.34. *For all integers $K \geqslant 2$, there exists $c_K > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2, 3, \dots K\}^n$, every vertex of $\vec{L}(n, \mathbf{k})$ can either be reached from at most*

$$N' = \frac{\log(K)}{\log(K) - \log(K-1)}$$

vertices or from every non-sink, with failure probability at most $e^{-c_K n}$.

Proof. Note first that it is sufficient to show that the result holds when n is large relative to K, since this covers all but finitely many cases for each K, and c_K can be chosen to handle these.

By the $\varepsilon = 1$ case of Theorem 3.8, there exists $c > 0$ such that if n is large relative to K, then for all $\mathbf{k} \in \{2, \dots, K\}^n$, the failure probability of the event that every vertex of $\vec{L}(n, \mathbf{k})$ can either be reached from at most $2K\log(K)$ vertices, or from every non-sink, is at most e^{-cn}. Hence, to prove the theorem it is enough to show that for all K there exists $c'_K > 0$ such that if n is large enough relative to K, then for all $\mathbf{k} \in \{2, \dots, K\}^n$, with failure probability at most $e^{-c'_K n}$, no vertices of $\vec{L}(n, \mathbf{k})$ can be reached from more than N' vertices but at most $2K\log(K)$ vertices.

This can be achieved by modifying the proof of Lemma 3.22. Defining Y_m as in Section 3.3.1, we need to show that $Y_m = 0$ for all $N' < m \leqslant 2K\log(K)$ with failure probability $e^{-c'_K n}$. Fix such an m, then as in the proof of Lemma 3.22 we have

$$\mathbb{E}[Y_m] \leqslant \left[4nK^2 \cdot \left(K^{1/m}\left(1 - \frac{1}{K}\right)\right)^{n-m}\right]^m.$$

In place of (3.1), it is not difficult to check that $m > N'$ ensures that $K^{1/m}(1-1/K) < 1 - \eta_K$ for some $\eta_K \in (0,1)$. Thus, for n large in terms of K (uniformly in m), we have

$$\mathbb{E}[Y_m] \leqslant \left(4nK^2(1-\eta_K)^{n-m}\right)^m \leqslant \left(1 - \eta_K^2\right)^{m(n-m)} \leqslant e^{-c''_K n},$$

for some $c''_K > 0$. Then, if n is large relative to K,

$$\sum_{N' < m \leqslant 2K\log(K)} \mathbb{E}[Y_m] \leqslant 2K\log(K) \cdot e^{-c''_K n} \leqslant e^{-c'_K n},$$

for some $c'_K > 0$, so by Lemma 3.16 we have that $Y_m = 0$ for all $N' < m \leqslant 2K\log(K)$ with failure probability at most $e^{-c'_K n}$, as required. □

Together, Theorem 3.33 and Theorem 3.34 essentially determine the 'correct' value for N as a function of K. Indeed, if we ignore the fact that r must be an integer in Theorem 3.33, then that theorem implies that when n is much larger than K, $\vec{L}(n, \mathbf{K})$ typically contains a vertex that can be reached from 'exactly' $\log(K)/(\log(K) - \log(K-1))$ vertices. Meanwhile Theorem 3.34 implies that typically every vertex which can be reached from more than this many vertices can be reached from every non-sink.

Although the improvement to the value of N represented by Theorem 3.34 is modest, it has the following consequence for $\vec{L}(n, \mathbf{2})$ and $\vec{L}(n, \mathbf{3})$ which may be of independent interest. Here, for $k \geqslant 2$ we write $\mathbf{k}$ for the all k's vector of the appropriate length.

Corollary 3.11. *With failure probability $e^{-\Omega(n)}$, every non-sink in $\vec{L}(n, \mathbf{2})$ can reach every non-source. The same is true for $\vec{L}(n, \mathbf{3})$. Conversely, for each $k \geqslant 4$, the probability that there is a non-sink in $\vec{L}(n, \mathbf{k})$ which cannot reach every non-source tends to 1 as $n \to \infty$.*

The positive direction of the corollary follows straightforwardly from Theorem 3.34 and the observation that every non-source in $\vec{L}(n, \mathbf{k})$ can be reached from at least $\min_i(k_i)$ vertices, and the negative direction follows immediately from setting $r = 1$ in Theorem 3.33. While the $\mathbf{k} = \mathbf{2}$ case of the corollary follows from Theorem 3.8, the $\mathbf{k} = \mathbf{3}$ case does not.

3.5 Proof of Proposition 3.12

In this section we give the proof of Proposition 3.12, which we restate below for convenience. The proof below actually yields a slightly better upper bound on the probability that $\vec{L}(n, \mathbf{k})$ is acyclic, but for clarity we have not included this in the statement.

Proposition 3.12. *There exists $c > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2, 3, \dots\}^n$, the probability that $\vec{L}(n, \mathbf{k})$ is acyclic is at most $\exp(-cnk^{n-2})$, where $k \coloneqq \min_i(k_i)$.*

For distinct $i, j \in [n]$, we define an *$\{i, j\}$-plane* of $V(n, \mathbf{k})$ to be a subset of $V(n, \mathbf{k})$ of size $k_i k_j$ whose elements pairwise differ in at most their ith and jth coordinates.

A subset of $V(n, \mathbf{k})$ will be called a *plane* of $V(n, \mathbf{k})$ if it is an $\{i, j\}$-plane for some i and j.

Proof of Proposition 3.12. We begin with the following claim.

Claim 3.35. *Let $k_1, k_2 \geqslant 2$ be integers, then $\vec{L}(2, (k_1, k_2))$ contains a cycle with probability at least $1/8$.*

Proof. Define a random process $\mathbf{X} = (X_0, X_1, X_2, \dots)$ coupled to $\vec{L}(2, (k_1, k_2))$ as follows. Let $X_0 = (1, 1)$ and for each $t \geqslant 1$, given $X_{t-1} \in [k_1] \times [k_2]$, if t is odd, let X_t be the winner of the line in coordinate 1 which contains X_{t-1}. If t is even, let X_t be the winner of the line in coordinate 2 which contains X_{t-1}. Thus, $\mathbf{X}$ is a random walk on $[k_1] \times [k_2]$ starting at $(1, 1)$, which at odd time steps traverses the available edge of $\vec{L}(2, (k_1, k_2))$ in the first coordinate direction (if there is one), and at even time steps traverses the available edge in the second.

Let T be the least time t at which there exists $1 \leqslant i < t$ such that X_i and X_t have the same first coordinate if t is odd, or the same second coordinate if t is even. It is not difficult to check (bearing in mind that we do not explore from X_0 in the second dimension at the start of the process) that for each $t < T$, when we explore from X_t we do so in an unexplored line.

Once we reach X_T, we already know the winner in the line we want to explore: it is the X_i with $1 \leqslant i < T$ that has the same first or second coordinate (depending on the parity of T) as X_T. Hence, $X_{T+1} = X_i$ for this i and the process is deterministic from here, with $X_{T+2} = X_{i+1}$ and so on. The process either becomes stationary at this point or enters a (non-trivial) cycle. It is straightforward to see that the process is stationary exactly when $X_T = X_{T-1}$, so if $X_T \neq X_{T-1}$, then $\vec{L}(2, (k_1, k_2))$ contains a cycle. Hence, let A be the event that $X_T \neq X_{T-1}$. We will show that $\mathbb{P}(A) \geqslant 1/8$.

Given T and X_{T-1}, X_T is chosen uniformly at random from among those vertices in the unexplored line through X_{T-1} that also belong to a previously explored line. There are $\lfloor T/2 \rfloor$ such vertices to choose from. Writing $\tau = \tau(k_1, k_2)$ for the maximum possible value of T, it follows that, for each $t \in \{2, \dots, \tau\}$, we have $\mathbb{P}(A \,|\, T = t) = (\lfloor t/2 \rfloor - 1)/\lfloor t/2 \rfloor$. This is at least $1/2$ for $t \geqslant 4$, so

$$\mathbb{P}(A) = \sum_{t=2}^{\tau} \mathbb{P}(A \,|\, T = t)\mathbb{P}(T = t) \geqslant \frac{\mathbb{P}(T \geqslant 4)}{2}.$$

We have $T \geqslant 4$ exactly when X_1 does not win its line in the second dimension and X_2 does not win its line in the first dimension. This event occurs with probability $(1 - 1/k_1)(1 - 1/k_2) \geqslant 1/4$, so $\mathbb{P}(A) \geqslant 1/8$, as required. □

Let $n \geqslant 2$ be an integer and let $\mathbf{k} \in \{2, 3, \dots\}^n$. By the claim, any given plane of $V(n, \mathbf{k})$ induces a cyclic subgraph of $\vec{L}(n, \mathbf{k})$ with probability at least $1/16$. In a family of planes which pairwise intersect in at most one vertex, each plane induces a cyclic subgraph of $\vec{L}(n, \mathbf{k})$ independently. The collection consisting of all $\{1, 2\}$-planes, all $\{3, 4\}$-planes, and so on, up to the $\{2\lfloor n/2 \rfloor - 1, 2\lfloor n/2 \rfloor\}$-planes, is such a family. For distinct $i, j \in [n]$, the number of $\{i, j\}$-planes in $V(n, \mathbf{k})$ is $\prod_{a \in [n] \setminus \{i,j\}} k_a$, so this family has size at least $\lfloor n/2 \rfloor \min(k_i)^{n-2}$, and the proposition follows. □

3.6 Adaptive dynamics and equilibrium selection

We now consider the implications of our results on game connectivity for games played over time according to adaptive dynamics. Some of what follows may appear slightly contrived from a mathematician's perspective, but the types of dynamics and notions of convergence to equilibria that we consider here are widely studied in economics (see, for example, [77, 86, 125, 155]). We begin by recalling some standard notions.

First, player i's *observation set* at time t, denoted o_i^t, is the set of information that i can observe at time t. Precisely what objects enter into this set varies depending on the regime under consideration, and below it will be made clear which regimes we are considering. For each integer $k \geqslant 2$, let O_k denote the set of all possible observation sets (under the given regime) for a player with action set $[k]$. A *strategy* for a player with action set $[k]$ is a function $f \colon O_k \to \Delta([k])$, where $\Delta([k])$ is the probability simplex over $[k]$. For $n \geqslant 2$ and $\mathbf{k} \in \{2, 3, \dots\}^n$, a *dynamic* on $\mathcal{G}(n, \mathbf{k})$ consists of a specification for what information enters into each player's observation set at each time, and a strategy f_i with action set $[k_i]$ for each player i.

The play of a game $g \in \mathcal{G}(n, \mathbf{k})$ under a given dynamic begins at time $t = 0$ at an initial action profile a^0 chosen arbitrarily. This informs each player's observation set o_i^1 according to the dynamic. At time $t = 1$, each player updates their action (randomly) according to $f_i(o_i^1)$, and we denote the new (random) action profile by a^1. The play continues in this manner, with each player updating their action at $t = 2$ according to $f_i(o_i^2)$ to produce an action profile a^2, and so on. We now define various types of dynamic.

Definition 3.36. A dynamic is *uncoupled* if at each time t, each player i's observation set contains (at most) their own preference relation $\succsim_i$ and the ordered history of play $a^0, \ldots, a^{t-1}$.

Definition 3.37. For an integer $m \geqslant 1$, an uncoupled dynamic is *m-recall* if at each time t, each player i's observation set contains (at most) the current time t, their own preference relation $\succsim_i$, and the ordered history of play $a^{t-m}, \ldots, a^{t-1}$ for the past m steps, or the full history of play if $t < m$.

Definition 3.38. An uncoupled and m-recall dynamic is *stationary* if at each time t each player i's observation set consists of their own preference relation $\succsim_i$ and the ordered history of play $a^{t-m}, \ldots, a^{t-1}$ for the past m steps, or the full history of play if $t < m$. Crucially, for $t \geqslant m$ the only information about the current time t available to the players is that $t \geqslant m$, so their strategies become time-independent after this point.

In what follows, we consider the following strong notion of convergence to a pure Nash equilibrium.[7]

Definition 3.39. A dynamic on $\mathcal{G}(n, \mathbf{k})$ *converges almost surely to a pure Nash equilibrium* of a game $g \in \mathcal{G}(n, \mathbf{k})$ if when g is played according to the dynamic from any initial action profile, almost surely there exists $T < \infty$ and a pure Nash equilibrium a^* of g such that $a^t = a^*$ for all $t \geqslant T$.

The following impossibility result is due to Hart and Mas-Colell [77] and Jaggard, Lutz, Schapira, and Wright [86].

Theorem 3.40 ([77, 86])**.** *For all $n \geqslant 3$ and $\mathbf{k} \in \mathbb{N}^n$ with $k_i \geqslant 2$ for all i (or $k_i \geqslant 3$ for all i if $n = 3$), there is no stationary, 1-recall, uncoupled dynamic on $\mathcal{G}(n, \mathbf{k})$ for which play converges almost surely to a pure Nash equilibrium in every generic game in the class that has one.*

However, we show below that as n grows (much more quickly than $\max_i(k_i)$), there are dynamics on $\mathcal{G}(n, \mathbf{k})$ of this type that converge almost surely to a pure Nash

[7] Almost sure convergence is a strong notion of convergence since it requires period-by-period play to eventually settle on, and never leave, a pure Nash equilibrium. We have focused on it, in part, because this is the notion that the impossibility results of [77, 86] pertain to. Weaker notions of convergence (see [153] for examples) admit different possibility results. For example, Young [155] showed that so-called 'trial-and-error learning' is an uncoupled (in fact, completely uncoupled) dynamic that, for any $\varepsilon > 0$, is at a pure Nash equilibrium for a $1 - \varepsilon$ proportion of time steps in any generic game that has one. Trial-and-error learning requires constant experimentation so there is always a positive probability of leaving a pure Nash equilibrium and wandering before settling on one again; such a dynamic is clearly not well suited to almost sure convergence.

equilibrium on all but a vanishingly small fraction of generic games in the class that have one.

Consider the following family of stationary, 1-recall, uncoupled dynamics. This dynamic is well known, with versions of it appearing in, for example, [140, 155].

Definition 3.41. Under the *best-response dynamic with inertia*, at each step t, independently of the other players, each player i sets a_i^t to be a best-response to a_{-i}^{t-1} with some fixed probability $p_i \in (0,1)$ and sets $a_i^t = a_i^{t-1}$ with complementary probability $1 - p_i$.[8]

Young [153] showed that, for any choice of parameters $p_i \in (0,1)$, this dynamic converges almost surely to a pure Nash equilibrium in every weakly acyclic game. We include a proof of this result in the case of generic weakly acyclic games in order to shed light on the link between the connectivity properties of games and the convergence of dynamics.

Theorem 3.42 ([153])**.** *For any choice of parameters $p_i \in (0,1)$, the best-response dynamic with inertia converges almost surely to a pure Nash equilibrium in every generic weakly acyclic game.*

Proof. Let g be a generic weakly acyclic game and let A be its set of action profiles. Fix some $a \in A$ which is not a pure Nash equilibrium, let t be an arbitrary time, and condition on the event that $a^t = a$. The vertex corresponding to a in the best-response graph of g is not a sink, so let i be a coordinate direction in which it has an outgoing edge. Then there exists some $\varepsilon > 0$ which is independent of a, t and i, such that with probability at least ε, at time $t+1$ player i changes their action by playing the best-response to a_{-i}, while all other players j repeat their existing action. In other words, after conditioning on $a^t = a$ for some non-sink a, the probability that at time step $t+1$ we move along any given out-edge of a in the best-response graph is at least ε.

Since g is weakly acyclic, for each $a \in A$ there is a path of length at most $|A|$ from a to a sink in the best-response graph. After fixing such a path for each a, we can repeatedly apply the above to lower bound the probability that at each step we move along the next edge in that path. Using the fact that the dynamic never leaves a pure Nash equilibrium once it arrives at one, this yields that for all times t and all

[8] The manner in which ties might be broken among multiple best-responses in non-generic games is immaterial for our purposes.

$a \in A$, conditioned on $a^t = a$, the probability that $a^{t+|A|}$ is a pure Nash equilibrium is at least $\varepsilon^{|A|}$. For each $m \in \{0, 1, 2, \dots\}$, denote by B_m the event that $a^{m|A|}$ is not a sink. It follows from the above that $\mathbb{P}(B_m \mid B_{m-1}) \leqslant 1 - \varepsilon^{|A|}$ for each $m \geqslant 1$, and so an inductive argument gives $\mathbb{P}(B_m) \leqslant (1-\varepsilon^{|A|})^m$. If the dynamic does not eventually settle at a pure Nash equilibrium, then B_m occurs for all m, but this has probability 0 since $(1 - \varepsilon^{|A|})^m \to 0$ as $m \to \infty$. This completes the proof of the theorem. □

Implications for adaptive dynamics Since every connected game is weakly acyclic, combining Theorem 3.5 and Theorem 3.42 yields the following.

Corollary 3.43. *There exists $c > 0$ such that for all integers $n \geqslant 2$ and all $\mathbf{k} \in \{2, 3, \dots\}^n$, if n is sufficiently large relative to $\max_i(k_i)$, then the proportion of games in*

$$\{g \in \mathcal{G}(n, \mathbf{k}) \colon g \text{ is generic and has a pure Nash equilibrium}\}$$

for which the best-response dynamic with inertia converges almost surely to a pure Nash equilibrium is at least $1 - e^{-cn}$.

In contrast to Theorem 3.40, this result shows that there is a stationary, 1-recall, uncoupled dynamic that converges almost surely to a pure Nash equilibrium in almost every large generic game that has one.

Of course, more is true. Since connected games are weakly acyclic, it follows from Theorem 3.5 that all existing results on dynamics converging in weakly acyclic games also apply in almost every large generic game that has a pure Nash equilibrium. We have focused here on the best-response dynamic with inertia but there are many more dynamics to which this applies. For example, it straightforwardly holds for the better-response dynamic with inertia and for a one-at-a-time version of the best-response dynamic[9] in which, at each step t, exactly one player i is selected at random from among all players to update their action, and this player plays a best-response to a_{-i}^{t-1}.

It is also known that 'adaptive play' converges almost surely to a pure Nash equilibrium in all weakly acyclic games [152], and the authors of [102] and [103] respectively describe regret-based and payoff-based dynamics that lead to play that is at a pure Nash equilibrium 'most of the time' in every weakly acyclic game. By Theorem 3.5, all of these results apply to almost every large generic game that has a pure Nash equilibrium, even when the notions of convergence are different from almost sure convergence of period-by-period play.

[9] Convergence properties of this one-at-a-time version were investigated in [81] via simulation.

Implications for equilibrium selection The implications of our work for adaptive dynamics discussed above do not require connectedness – weak acyclicity suffices. Nevertheless, connected games have the following important dynamical property: suppose that the best-response dynamic with inertia has reached a pure Nash equilibrium but is then shocked out of this equilibrium by, for example, some player committing an error and playing a non-best-response action. If the game is connected, then every non-sink can reach *every* sink, so the shocked dynamic can go on to converge to *any* pure Nash equilibrium via best-response paths.[10] This is not necessarily the case in weakly acyclic games.

This property of connected games has very recently been exploited by Newton and Sawa [115], who have applied our results in the context of evolutionary game theory. We now briefly outline their findings.

For games with multiple pure Nash equilibria, evolutionary game theorists seek to understand at which of the equilibria simple adaptive dynamics will spend most of their time. For example, consider a perturbed version of the best-response dynamic with inertia in which, at each time, any updating player plays a best-response with probability $1 - \varepsilon$ for some $\varepsilon \in (0, 1)$, and otherwise selects an action uniformly at random. The Markov process induced by this dynamic is irreducible and positive recurrent, and hence has a unique invariant distribution, μ_ε. Moreover, it is not difficult to see that these invariant distributions have a pointwise limit, μ, as $\varepsilon \to 0$. The *stochastically stable* states of the perturbed best-response dynamic with inertia are the action profiles that are assigned positive probability under μ. Roughly, these are the states at which the dynamic spends "most" of its time. The definition of stochastic stability can, of course, be applied to other "perturbed" dynamics in the obvious way.

Commonly used methods for determining stochastic stability include the minimum-cost tree technique and the radius-coradius technique [44, 55, 89, 90, 152], but applying these can be complicated because they require checking global properties: a stochastically stable state must be hard to leave (requiring many 'errors' to exit) and easy to enter (requiring few 'errors' to arrive at from other points in the game).

Newton and Sawa [115] observe that, in connected games, because it is very easy to move from each non-equilibrium point to *any* equilibrium point, determining the

[10] This type of 'fragility' is investigated by Rudov [127] in the context of matching markets.

stochastically stable states reduces to checking only 'one-shot deviations' from equilibrium, which is a local property regarding the likelihood of exiting an equilibrium point. Applying our results, they conclude that in almost every large generic game that has a pure Nash equilibrium, stochastic stability is determined by this very simple 'one-shot' property. They go on to determine which pure Nash equilibria (according to their welfare properties) are selected by various evolutionary dynamics in typical large games.

To reiterate, these results on equilibrium selection require our strong notion of connectivity; they do not apply in weakly acyclic games that are not also connected.

3.7 Open questions

We have approached the problem of finding adaptive dynamics that converge to pure Nash equilibria by studying connectivity properties of games rather than studying the properties of the dynamics themselves. We hope to have demonstrated that this approach can deliver interesting conclusions, but many open questions still remain. We outline some of them below.

(i) We have worked in the regime in which the number of players is much larger than the maximum number of actions per player. One could alternatively consider the setting where the number of players is fixed (or growing slowly) and the number of actions gets large. Results in this regime are limited to two-player games. Heinrich, Jang, Mungo, Pangallo, Scott, Tarbush, and Wiese [81] showed that almost no generic two-player game that has a pure Nash equilibrium is weakly acyclic as the number of actions gets large for both players. On the other hand, Amiet, Collevecchio, and Hamza [6] showed that almost *every* generic two-player game that has a pure Nash equilibrium is *globally* weakly acyclic as the number of actions gets large (again, for both players). However there are, to our knowledge, no results where the number of players is fixed above two.

The 'large number of actions' regime (where n is fixed and k grows) behaves very differently from the 'large number of players' regime that we consider in this chapter. In the latter, each vertex of the best-response graph is incident to at least n edges, so n being large is likely to contribute to greater connectivity. When instead it is the number of actions that is large, vertices are incident to only a fixed number of lines, and this leads to the possibility of sinks with relatively

small 'basins of reachability'. Characterising this regime will require different arguments from those employed in our proof of Theorem 3.8.

(ii) The interest in uncoupled dynamics stems from the question of whether there are informationally undemanding dynamics that are guaranteed to lead to a pure Nash equilibrium whenever there is one. Completely uncoupled dynamics have even lower informational requirements than uncoupled ones: a dynamic is *completely uncoupled* if, at each time t, each player's observation set contains only their own realised utility payoffs and their own past actions. Babichenko [13] showed that there is no completely uncoupled dynamic for which the period-by-period play converges almost surely to a pure Nash equilibrium in every generic game that has one, and moreover that there exist obstructions with arbitrarily large numbers of players and numbers of actions per player. We have not considered completely uncoupled dynamics here since they rely intrinsically on games with utility functions, which are outside our scope. Nevertheless, further analysis of game connectivity properties, either those we have studied here or others, may yield positive results on completely uncoupled dynamics.

(iii) Since our focus is not on any specific dynamic, we have not addressed the question of the speed of convergence to equilibrium [10]. It is known that adaptive dynamics can take a very long time to converge [73], but greater knowledge of connectivity properties may help to establish general results on the speed of convergence in some classes of games.[11]

(iv) Amiet, Collevecchio, Scarsini, and Zhong [7] derive some connectivity properties of non-generic games with two actions per player but there are, to our knowledge, no results for non-generic games with more than two actions per player.

(v) Our analysis of convergence to pure Nash equilibrium did not consider the 'quality' of these equilibria. Pradelski and Young [118], for example, describe a completely uncoupled dynamic that spends most of its time at Pareto optimal equilibria.[12] Again, further analysis of game connectivity properties may help to more broadly address the question of convergence to efficient equilibria.

(vi) As noted in the points above, it might be fruitful to investigate the prevalence of games with other types of connectivity properties. For example, it may be worth

[11] Convergence can be fast in potential games (e.g. see [12]) and in anonymous games (e.g. see [14]).

[12] A pure Nash equilibrium is *Pareto optimal* if there are no other pure Nash equilibria at which at least one player is better off and no player is worse off.

considering graphs that correspond to deviations by multiple players at a time, rather than just one.

Chapter 4

Induced subgraphs of induced subgraphs of large chromatic number

Recall from Section 1.3 that a hereditary graph class $\mathcal{G}$ is said to be χ-*bounded* if there exists a function f such that $\chi(G) \leqslant f(\omega(G))$ for all $G \in \mathcal{G}$. An important technique for determining whether a class $\mathcal{G}$ is χ-bounded is to 'clean up' the class by showing that every graph in $\mathcal{G}$ with very large chromatic number has an induced subgraph that still has large chromatic number but is in some sense structurally simpler. It then suffices to analyse the subclass of these structurally simpler graphs. An extreme version of this approach was put forward independently by Galvin and Rödl (see [108]) and Esperet (see [136]), who conjectured that every graph with very large chromatic number contains either a large clique or an induced triangle-free subgraph with large chromatic number. If true, this would imply that whether or not a hereditary graph class is χ-bounded depends only on the triangle-free graphs in the class.

This conjecture, however, was recently disproved by Carbonero, Hompe, Moore, and Spirkl [35], who constructed K_4-free graphs of arbitrarily large chromatic number, all triangle-free induced subgraphs of which have chromatic number at most four. Shortly afterwards, Briański, Davies, and Walczak [30] generalised this result from triangles to cliques of prime order p, constructing K_{p+1}-free graphs with arbitrarily large chromatic number, all K_p-free induced subgraphs H of which satisfy $\chi(H) \leqslant \omega(H)^{\omega(H)^2}$. In an update [31] to their paper (appearing independently after the results in this chapter were first made available online) they improved the bound on $\chi(H)$ to $\binom{3\omega(H)+1}{3}$ and removed the condition that p is prime.

It follows that the very strong style of 'cleaning up' proposed by Galvin and Rödl and by Esperet is not viable: we cannot simply drop to a subclass with bounded clique

number. Nevertheless, one might still wonder whether *some* induced subgraph F can be excluded from the class. The main result of this chapter is that, in fact, this is not the case for any non-trivial F. Recall that we say a graph is *F-free* if it does not contain an induced copy of F.

Theorem 1.8. *For every graph F with at least one edge, there is a constant, c_F, for which there exist graphs G of arbitrarily large chromatic number and the same clique number as F such that every F-free induced subgraph of G is c_F-colourable.*

For triangle-free F, Theorem 1.8 guarantees a triangle-free G. This raises the question of whether G can always be taken to have the same girth as F. We conjecture that this is the case (see Section 4.7), but have only been able to prove the following weaker statement in which girth is replaced by odd girth.

Theorem 4.1. *For every non-bipartite graph F, there is a constant, c'_F, for which there exist graphs G of arbitrarily large chromatic number and the same odd girth as F such that every F-free induced subgraph of G is c'_F-colourable.*

Theorems 1.8 and 4.1 both easily extend to finite families of graphs. Indeed, if $\mathcal{F}$ is a finite family of graphs (at least one of which contains an edge) and $G_{\mathcal{F}}$ is the disjoint union of the members of $\mathcal{F}$, then applying Theorem 1.8 to $G_{\mathcal{F}}$ shows that there is some constant $c_{\mathcal{F}}$ and graphs G of arbitrarily large chromatic number and the same clique number as the largest clique number of a graph in $\mathcal{F}$ such that every induced subgraph of G with chromatic number greater than $c_{\mathcal{F}}$ contains every member of $\mathcal{F}$ as an induced subgraph. Does this phenomenon occur for infinite families? We will say that a family of graphs $\mathcal{F}$ is *durable* if there exists an infinite graph H with infinite chromatic number such that every infinite-chromatic induced subgraph of H contains every member of $\mathcal{F}$ as an induced subgraph. All finite families of graphs are durable but the situation for infinite families is more complicated. We obtain an effective test for deciding whether a family of unbounded chromatic number is durable which, for example, gives the following result.

Theorem 4.2. *For every integer g, the family of graphs with girth at least g is not durable.*

In the bounded chromatic number case we obtain a link to χ-boundedness, which in particular implies the following.

Theorem 4.3. *The family of paths is durable.*

We also extend and generalise Theorem 1.8 to the settings of tournaments and hypergraphs (see Sections 4.4 and 4.5 for definitions not given below). For tournaments we prove the following.

Theorem 4.4. *For every tournament T, there is a constant, C_T, for which there exist tournaments S of arbitrarily large chromatic number such that every T-free subtournament of S has chromatic number at most C_T.*

For hypergraphs, there are two notions of chromatic number: the usual one where no edge is monochromatic and the strong one where no two vertices in an edge have the same colour. Pleasingly, our generalisation of Theorem 1.8 uses the usual notion of chromatic number for the constructed hypergraph, but the strong notion for its induced subhypergraphs. The statement also has a very general analogue of a clique.

A hypergraph is said to be *strongly t-colourable* if its vertices can be t-coloured such that no edge contains two vertices of the same colour (or, equivalently, if the graph obtained by replacing each edge of the hypergraph by a clique has chromatic number at most t). A hypergraph *covers* a pair of vertices if it has an edge containing both of them. Theorem 4.5 applies to general (not necessarily uniform) hypergraphs, but we assume throughout this chapter that all edges have size at least two.

Theorem 4.5. *For every hypergraph $\mathcal{F}$ with at least one edge, there is a constant, $c_{\mathcal{F}}$, for which there exist hypergraphs $\mathcal{G}$ of arbitrarily large chromatic number such that every $\mathcal{F}$-free induced subhypergraph of $\mathcal{G}$ is strongly $c_{\mathcal{F}}$-colourable. Moreover, we can take $\mathcal{G}$ such that if $\mathcal{G}$ covers every pair from some $X \subseteq V(\mathcal{G})$, then $\mathcal{G}[X]$ is an induced subhypergraph of $\mathcal{F}$.*

It follows from the moreover part of the theorem that the edge sizes which appear in $\mathcal{G}$ are the same as those which appear in $\mathcal{F}$, and that for all k the maximum size of a k-uniform clique in $\mathcal{G}$ is the same as that in $\mathcal{F}$. In particular, this theorem contains Theorem 1.8.

Our constructions for Theorems 1.8, 4.1, and 4.5 all use a common framework: we start with a base structure of large chromatic number and add edges according to 'distances' in the base structure. Using generalised Sidon sets, we are able to choose these 'distances' to encode F or $\mathcal{F}$ and give precise control over where its copies appear.

Chapter outline In Section 4.1 we describe the base graphs used in the proofs of Theorems 1.8 and 4.1, then we complete the proofs of these theorems in Sections 4.2 and 4.3 respectively. Section 4.4 is concerned with tournaments and contains the proof of Theorem 4.4, then in Section 4.5 we turn our attention to hypergraphs and prove Theorem 4.5. Our results on the durability of infinite families of graphs are covered in Section 4.6, before we close with a discussion of some open problems in Section 4.7.

This chapter is based on joint work with António Girão, Freddie Illingworth, Emil Powierski, Alex Scott, Youri Tamitegama, and Jane Tan [62].

4.1 Orientations of classical constructions

Our constructions for Theorems 1.8 and 4.1 start with orientations of graphs with high chromatic number and restricted clique number or girth. Traditionally the underlying graphs in such orientations have been the well known Zykov graphs [158]. Orientations of these were first considered by Kierstead and Trotter [93], and then used in the constructions of both Carbonero, Hompe, Moore, and Spirkl [35] and Briański, Davies, and Walczak [30]. For the proof of Theorem 1.8, we require precisely the following properties of the underlying graphs X_n and their orientations $\overrightarrow{X_n}$:

(a) $\chi(X_n) = n$,

(b) $\overrightarrow{X_n}$ is acyclic, that is, it contains no directed cycle, and

(c) for all $u, v \in V(\overrightarrow{X_n})$, there is at most one directed path from u to v in $\overrightarrow{X_n}$.

The Zykov graphs and their orientations given by Kierstead and Trotter satisfy these conditions and so suffice for the proof of Theorem 1.8. In fact, any X_n with chromatic number n and girth greater than $2n-2$ has some orientation $\overrightarrow{X_n}$ satisfying these three properties (see Lemma 4.7).

For the proof of Theorem 4.1 we require one further property in order to rule out short odd cycles (in particular, we cannot take Zykov graphs as the base graphs in the proof of Theorem 4.1). Define a *change of direction* in a digraph to be a pair of edges incident to a common vertex that are both directed towards or both directed away from that vertex. For a given positive integer g we ask that $\overrightarrow{X_n}$ satisfies:

(d) for every cycle in X_n, the corresponding oriented cycle in $\overrightarrow{X_n}$ has at least g changes of direction.

Strictly speaking, $\overrightarrow{X_n}$ and condition (d) depend on the choice of positive integer g, but we suppress this dependence in the notation. Note that if $g \geqslant 3$, then property (d) implies both (b) and (c). Lemma 4.7 below states that any X_n with chromatic number n and girth greater than $(g-1)(n-1)$ where $g \geqslant 3$ has an orientation $\overrightarrow{X_n}$ satisfying all four properties.

By property (c), it is natural to view $V(\overrightarrow{X_n}) = V(X_n)$ as a poset with a strict partial order based on *reachability*: for distinct vertices u and v, we write $u < v$ if there is a directed path from u to v in $\overrightarrow{X_n}$. Note that there is a correspondence between chains in this poset and sets of vertices that lie on a single directed path in $\overrightarrow{X_n}$. For vertices $u < v$ we can define the *distance* between u and v, written $d(u, v)$, to be the length of the unique directed path from u to v. The uniqueness of directed paths ensures that $d(u, w) = d(u, v) + d(v, w)$ for vertices $u < v < w$.

We continue with the promised constructions. The following proposition, which is the acyclic case of the Gallai–Roy–Hasse–Vivater theorem (see, for example, [111, Theorem 3.13]), provides a strong relationship between orientations and chromatic number.

Proposition 4.6. *A graph L is k-colourable if and only if it has an acyclic orientation $\vec{L}$ which contains no directed paths of length k.*

Proof. First suppose that L is k-colourable and let $c\colon V(L) \to \{1, \dots, k\}$ be a proper k-colouring. For each edge $e = uv$ of L, if $c(u) < c(v)$, then orient e from u to v and if $c(v) < c(u)$, then orient e from v to u. For any directed path $u_1u_2 \dots u_t$ in the resulting orientation $\vec{L}$, we have $1 \leqslant c(u_1) < c(u_2) < \cdots < c(u_t) \leqslant k$ and so the orientation is acyclic and the path has length at most $k - 1$.

On the other hand, if the acyclic orientation $\vec{L}$ has no directed paths of length k, then we may simply colour each vertex v of L by the length of the longest directed path in $\vec{L}$ ending at v. □

Erdős [47] showed that for any integers $g, n \geqslant 2$ there are graphs of chromatic number n and girth at least g. This together with the following lemma (which was pointed out to us by Nešetřil and replaces a careful analysis of a construction of his and Rödl's [114] that we previously used to the same effect) shows that there are indeed digraphs with properties (a) to (d).

Lemma 4.7. *Let $g \geqslant 3$ and $n \geqslant 2$ be integers. If a graph X_n has chromatic number n and girth greater than $(g-1)(n-1)$, then X_n has an orientation $\overrightarrow{X_n}$ satisfying properties* (a), (b), (c), *and* (d).

Proof. The graph X_n has chromatic number n so, by Proposition 4.6, has an acyclic orientation $\overrightarrow{X_n}$ which contains no directed paths of length n. In particular, $\overrightarrow{X_n}$ automatically satisfies properties (a) and (b). Consider a cycle C in X_n. The length of C is greater than $(g-1)(n-1)$ and every directed path in $\overrightarrow{X_n}$ has length at most $n-1$, so C has at least g changes of direction in $\overrightarrow{X_n}$. Hence property (d) is satisfied. Finally, as $g \geqslant 3$, this implies that $\overrightarrow{X_n}$ satisfies property (c) also. □

4.2 Clique number

In this section, we will give the proof of Theorem 1.8. We have not tried hard to optimise the bounds; however, writing $f = |V(F)|$, the proof gives

$$\begin{aligned} c_F &\leqslant (32f^4)^{|E(F)|}, \\ c_F &\leqslant \big((2+o(1))f^4\big)^{|E(F)|}. \end{aligned} \tag{4.1}$$

Using results from the updated version of Briański, Davies, and Walczak's paper [31], this can be improved to $c_F = O(f^9)$ by noting that the graph G we construct in Section 4.2.2 can be taken to be a subgraph of their $G_{n,p}$. We begin by sketching the proof strategy, and then discuss B_3-sets in Section 4.2.1, before giving the details of the proof in Section 4.2.2.

The constructions in [30, 35, 93] start with oriented (Zykov) graphs $(\overrightarrow{X_n})$ satisfying properties (a) through (d), and add new edges $\overrightarrow{uv}$ (or sometimes $\overrightarrow{vu}$ in [35]) between vertices $u < v$ whenever p does not divide $d(u,v)$ for some prime p. These papers use arbitrary primes p, $p = 3$, and $p = 2$ respectively. Our construction is similar in that, for some prime p, we add edges to $\overrightarrow{X_n}$ based on the residues modulo p of the distances between endpoints. However, the set of residues for which we add edges is now more sophisticated. Let $\vec{G}$ be the digraph so obtained, and let G be its underlying undirected graph. As X_n is a subgraph of G, we know that $\chi(G) \geqslant n$. The extra edges will be added so that the following properties hold:

- $\omega(G) = \omega(F)$,
- $\vec{G}$ is acyclic, and

- if there is a long (depending only on F) directed path in $\vec{G}$ all of whose edges correspond to the same distance modulo p, then some vertices of this path induce a copy of F in G.

The second and third properties are particularly useful in light of Proposition 4.6.

Now taking G as above, suppose that H is an induced subgraph of G which does not contain F as an induced subgraph. Let H_i be the subgraph of H which consists of those edges corresponding to distance i modulo p. By the third bullet point, H_i cannot contain a long directed path. Thus, Proposition 4.6 and the second bullet point imply that H_i has chromatic number bounded in terms of F. It will turn out that $H_1 = X_n$, so taking a product colouring (over the possible i) will then show that H itself must have chromatic number at most some c_F.

4.2.1 B_3-sets

To guarantee that G has the same clique number as F our construction will utilise B_3-sets. A set of integers $S = \{a_1 < a_2 < \cdots < a_k\}$ is a *B_3-set* if the sums

$$a_{i_1} + a_{i_2} + a_{i_3}, \qquad 1 \leqslant i_1 \leqslant i_2 \leqslant i_3 \leqslant k$$

are all different. A simple example of a B_3-set is the set of powers of 4 — it would suffice to use an initial segment of this for our arguments, although the resulting bound on c_F would be much worse. B_3-sets (and more generally B_h-sets, which we discuss in Section 4.3.1) were introduced by Bose and Chowla [28] as a generalisation of Sidon sets [137], which are precisely the B_2-sets (i.e. sets of integers where all pairwise differences are distinct). In particular, every B_3-set is also a B_2-set. Let S be a B_3-set and let

$$D = \{s - s' \colon s' < s \text{ and } s, s' \in S\}$$

be the *difference set of S*. The sets S and D satisfy the following two claims which motivate our interest in B_3-sets.

Claim 4.8 (Triangle fact)**.** *Let S be a B_3-set and let D be its difference set. Suppose that D contains not necessarily distinct d_1, d_2 such that $d_1 + d_2 \in D$. Then there are $b_1 < b_2 < b_3$ all in S such that $\{b_2 - b_1, b_3 - b_2\} = \{d_1, d_2\}$.*

Proof. Suppose that $d_1, d_2, d_1 + d_2$ are all in D. Then there are $x_1, x_2, x_3, y_1, y_2, y_3$ all

in S with $d_1 = y_1 - x_1$, $d_2 = y_2 - x_2$, and $d_1 + d_2 = y_3 - x_3$. Hence,

$$\begin{aligned}(y_1 - x_1) + (y_2 - x_2) &= y_3 - x_3,\\ \Rightarrow y_1 + y_2 + x_3 &= x_1 + x_2 + y_3.\end{aligned}$$

But S is a B_3-set, so y_1, y_2, x_3 must be x_1, x_2, y_3 in some order. The elements of D are positive integers so $x_i \neq y_i$ for all i. Hence y_1 is either x_2 or y_3.

- If $y_1 = x_2$, then $b_1 = x_1$, $b_2 = y_1 = x_2$, and $b_3 = y_2$ satisfy $(b_2 - b_1, b_3 - b_2) = (d_1, d_2)$.
- If $y_1 = y_3$, then $y_2 = x_1$ (as $y_2 \neq x_2$) and so $b_1 = x_2$, $b_2 = x_1 = y_2$, and $b_3 = y_1$ satisfy $(b_2 - b_1, b_3 - b_2) = (d_2, d_1)$. □

The preceding claim can be leveraged to cover more distances.

Claim 4.9 (Clique fact)**.** *Let S be a B_3-set and let D be its difference set. Let ℓ be a natural number and suppose that D contains not necessarily distinct $d_1, \ldots, d_\ell$ such that every sum of the form*

$$\sum_{i_1 \leqslant j \leqslant i_2} d_j \tag{4.2}$$

with $1 \leqslant i_1 \leqslant i_2 \leqslant \ell$ is in D. Then there are $b_1 < \cdots < b_{\ell+1}$ all in S such that either $b_{i+1} - b_i = d_i$ for all i, or $b_{i+1} - b_i = d_{\ell+1-i}$ for all i.

Proof. We proceed by induction on ℓ. The case $\ell = 1$ is trivial and the case $\ell = 2$ is the triangle fact, so suppose $\ell \geqslant 3$. By the induction hypothesis, there are $b_1 < \cdots < b_\ell$ such that either $b_{i+1} - b_i = d_i$ for all $i < \ell$, or $b_{i+1} - b_i = d_{\ell-i}$ for all $i < \ell$. We may assume that $b_{i+1} - b_i = d_i$ for all $i < \ell$, as the other case follows by a symmetric argument.

Let us say that a difference $d \in D$ is associated with $b \in S$ if there is some $c \in S$ such that $|b - c| = d$. As S is a Sidon set, $d_{\ell-1}$ is only associated with $b_{\ell-1}$ and b_ℓ. Thus, applying the triangle fact to $d_{\ell-1}$ and d_ℓ we find that there is $b \in S$ such that either $b_{\ell-1} - b = d_\ell$ or $b - b_\ell = d_\ell$. In the latter case we are done, so assume that $b_{\ell-1} - b = d_\ell$.

Similarly, $d \coloneqq \sum_{j=1}^{\ell-1} d_j$ is only associated with b_1 and b_ℓ, so by applying the triangle fact to d and d_ℓ we find that there is some $b' \in S$ such that either $b_1 - b' = d_\ell$ or $b' - b_\ell = d_\ell$. In the latter case we are done so we assume that $b_1 - b' = d_\ell$. But then $b' < b_1 < b_{\ell-1}$ are all associated with d_ℓ while S is a Sidon set. This is a contradiction, as required. □

The existence of large B_3-sets is guaranteed by the following theorem of Bose and Chowla [28, Theorem 2]. As in previous chapters, we use the standard shorthand $[n] = \{1, \dots, n\}$.

Theorem 4.10 ([28])**.** *Let q be one more than a prime power. There is a B_3-subset of $[q^3]$ of size q that contains both 1 and 2.*

Given a natural number f, using the well known fact that the gap between consecutive primes $q < q'$ is $o(q)$ (see, for instance, [15] for an up-to-date bound), Theorem 4.10 implies that for $N = (1 + o(1))f^3$, there is a B_3-subset of $[N]$ of size at least f that contains both 1 and 2. Bertrand's postulate gives a non-asymptotic bound: the smallest prime $p = q - 1$ greater than $f - 1$ is less than $2f - 1$ and so Theorem 4.10 gives a B_3-subset of $[q^3] \subseteq [8f^3]$ of size at least f that contains both 1 and 2.

4.2.2 Proof of Theorem 1.8

Fix a graph F with at least one edge and let $f = |V(F)|$. Let p be a prime number chosen to be large enough that $\big[(p-1)/2\big]$ has a B_3-subset S of size f that contains 1 and 2. By the discussion at the end of Section 4.2.1, it suffices to have $(p-1)/2 \geqslant \min\big\{(1+o(1))f^3, 8f^3\big\}$. Using the result on prime gaps and Bertrand's postulate we may choose such a p with $p \leqslant \min\big\{(2+o(1))f^3, 32f^3\big\}$.

Let the elements of S be $s_1 = 1, s_2 = 2, s_3, \dots, s_f$ in ascending order and let the difference set of S be $D = \{s_j - s_i \colon i < j\}$. As F has at least one edge we may take a copy F^* of F with vertex set S which includes the edge $s_1 s_2$. We write

$$E = \{s_j - s_i \colon i < j \text{ and } s_i s_j \in E(F^*)\}$$

for the set of distances corresponding to edges of F^*. Since $s_1 = 1$ and $s_2 = 2$, we have $1 \in E$. Note that $E \subseteq D$. Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = E + p\mathbb{Z}$.

For a fixed large integer n, let $\overrightarrow{X_n}$ and X_n be (di)graphs satisfying conditions (a), (b), and (c) from Section 4.1, which were shown to exist in that section. Recall that for vertices $u < v$ (in the reachability ordering) of these graphs we write $d(u, v)$ for the length of the unique directed path from u to v in $\overrightarrow{X_n}$. We now define a digraph $\vec{G}$ with vertex set $V(\overrightarrow{X_n})$ as follows.

$$\text{The edge } \overrightarrow{uv} \text{ is present if } u < v \text{ and } d(u, v) \in \overline{E}.$$

The digraph $\vec{G}$ is acyclic, since the presence of the edge $\overrightarrow{uv}$ implies that $u < v$ in the reachability order. Let G be the underlying undirected graph of $\vec{G}$.

We now verify that G satisfies the conditions of Theorem 1.8. Since $1 \in E$, the graph X_n is a subgraph of G and so $\chi(G) \geqslant n$. Next we show that $\omega(G) \leqslant \omega(F)$. To this end, suppose that vertices $v_1, \ldots, v_{\ell+1}$ form a clique in G. Since any two of these vertices are $<$-comparable, we may assume that $v_1 < \cdots < v_{\ell+1}$. For $i \in [\ell]$ let $d_i \in \{0, \ldots, p-1\}$ be the residue modulo p of $d(v_i, v_{i+1})$, and note that

$$d(v_{i_1}, v_{i_2+1}) \equiv \sum_{i_1 \leqslant j \leqslant i_2} d_j \bmod p$$

for all $1 \leqslant i_1 \leqslant i_2 \leqslant \ell$. Since $v_1, \ldots, v_{\ell+1}$ form a clique in G, all these distances $d(v_{i_1}, v_{i_2+1})$, and hence all sums of the form in (4.2), are in $\overline{E} \subseteq \overline{D}$. We claim that these sums are actually in D. It is clear that $d_j \in D$ for all j, and since $D \subseteq [(p-1)/2]$ we have that the two-term sums of the form in (4.2) are in D as well. Iterating this shows that all sums of the form in (4.2) are in D.

It now follows from the clique fact that there are $b_1 < \cdots < b_{\ell+1}$ all in S such that either $b_{i+1} - b_i = d_i$ for all i, or $b_{i+1} - b_i = d_{\ell+1-i}$ for all i. All pairwise distances between elements of $B = \{b_1, \ldots, b_{\ell+1}\}$ are of the form in (4.2) and so are in $\overline{E}$. But since $B \subseteq S$, they must all actually be in E. By the definition of E and the fact that S is a Sidon set, the vertices $b_1, \ldots, b_{\ell+1}$ induce a clique in F^* and so $\ell + 1 \leqslant \omega(F)$, as required.

We now turn to the last required property of G, that all of its induced subgraphs of chromatic number greater than some c_F contain an induced copy of F. Start by colouring each edge uv of G (and $\vec{G}$), where $u < v$, with the colour $d(u, v) \bmod p$. Suppose that $\vec{P}$ is a directed path in $\vec{G}$ of length pf which is monochromatic in colour i for some $i \in E$. Label the first vertex of this path by v_1. Recall that the elements of S are $s_1 < \cdots < s_f$. Let v_2 be the first vertex on path $\vec{P}$ such that $d(v_1, v_2) \equiv s_2 - s_1 \bmod p$. Since p is prime and $i \neq 0$, there are at most $p-1$ edges of $\vec{P}$ between v_1 and v_2 on this path. Next, define v_3 to be the first vertex of $\vec{P}$ such that $d(v_2, v_3) \equiv s_3 - s_2 \bmod p$, and continue in this manner to define $v_4, \ldots, v_f$. Note that for each i there are at most $p-1$ edges between v_i and v_{i+1} on path $\vec{P}$, so the path is long enough that these vertices can be found.

The way in which we have chosen the vertices v_i means that, for each pair $i < j$, we have $d(v_i, v_j) \equiv s_j - s_i \bmod p$. Then, by the construction of G and the fact that S is a Sidon set, the edge $v_i v_j$ is present in the graph exactly when there is an edge between s_i and s_j in F^*. It follows that G contains an induced copy of F with vertex set $\{v_1, \ldots, v_f\}$.

Now let H be an induced subgraph of G which does not contain an induced copy of F, and let $\vec{H}$ be the subdigraph of $\vec{G}$ induced by the same set of vertices. For each $i \in E$, let H_i be the subgraph of H consisting of the i-coloured edges and define $\vec{H_i}$ similarly. By the above, $\vec{H_i}$ contains no directed paths of length pf and hence H_i is pf-colourable by Proposition 4.6. Take such a proper colouring $\chi_i \colon V(H_i) \to [pf]$ for each i. Then the product colouring $\chi \colon V(H) \to [pf]^{|E|}$ given by $\chi(v) = (\chi_i(v) \colon i \in E)$ is a proper colouring of H. Hence, we may take c_F to be $(pf)^{|E|}$. Finally, we note that $|E| = |E(F)|$ and recall that $p \leqslant \min\{(2+o(1))f^3, 32f^3\}$ to obtain the bounds given by (4.1).

4.3 Odd girth

The proof of Theorem 4.1 shares some ingredients with that of Theorem 1.8. In this exposition, we will pay particular attention to the differences.

4.3.1 B_h-sets

Given a natural number $h \geqslant 2$, we can define the notion of B_h-sets analogously to B_3-sets. That is, a set of integers $S = \{a_1 < a_2 < \cdots < a_k\}$ is a *B_h-set* if the sums

$$a_{i_1} + a_{i_2} + \cdots + a_{i_h}, \qquad 1 \leqslant i_1 \leqslant i_2 \leqslant \cdots \leqslant i_h \leqslant k$$

are all different. Note that a B_h-set is also a $B_{h'}$-set for any $h' < h$, and in particular is a Sidon set. The powers of h form a B_h-set but Bose and Chowla [28, Theorem 2] provide a more efficient construction.

Theorem 4.11 ([28])**.** *Let h be a positive integer and q be one more than a prime power. There is a B_h-subset of $[q^h]$ of size q that contains both 1 and 2.*

The importance of B_h-sets for our proof of Theorem 4.1 is the following claim. We will only make use of the 'moreover' part but we need the main statement to apply induction. By a *circuit*, we mean a closed walk in which vertices may be repeated but edges may not.

Claim 4.12 (Cycle fact)**.** *Let $2 \leqslant \ell \leqslant h$ be integers, let S be a B_h-set, and let $D = \{s - s' \colon s' < s \text{ and } s, s' \in S\}$ be the difference set of S. Suppose that D contains not necessarily distinct $d_1, \ldots, d_\ell$ such that for some $1 \leqslant s \leqslant \ell - 1$ we have*

$$d_1 + \cdots + d_s = d_{s+1} + \cdots + d_\ell. \tag{4.3}$$

Let M be the multigraph on vertex set S built by adding, for each $1 \leqslant i \leqslant \ell$, an edge e_i between the pair of vertices whose difference is d_i. Then the edge set of M can be decomposed into a collection of circuits. Moreover, if ℓ is odd, then the graph obtained from M by deleting duplicate edges contains an odd cycle.

Proof. We prove that the edge set of M can be decomposed into a collection of circuits by induction on ℓ. If $\ell = 2$, then we have $d_1 = d_2$ and M is just a pair of edges between the same two vertices.

Suppose that $\ell \geqslant 3$. Let $x_1, \dots, x_\ell, y_1, \dots, y_\ell \in S$ be such that $d_i = y_i - x_i$ for all $1 \leqslant i \leqslant \ell$. From equation (4.3) we obtain

$$y_1 + \cdots + y_s + x_{s+1} + \cdots + x_\ell = x_1 + \cdots + x_s + y_{s+1} + \cdots + y_\ell$$

and so, as S is a B_h-set and $\ell \leqslant h$, we have that $y_1, \dots, y_s, x_{s+1}, \dots, x_\ell$ must be $x_1, \dots, x_s, y_{s+1}, \dots, y_\ell$ in some order. In particular, $y_1 = x_i$ for some $1 \leqslant i \leqslant s$ or $y_1 = y_j$ for some $s + 1 \leqslant j \leqslant \ell$. In the first case we may assume by relabelling that $y_1 = x_2$ ($y_1 \neq x_1$ as $0 \notin D$). Then $d \coloneqq d_1 + d_2 = y_2 - x_1$, which is in D as $y_2 > x_2 = y_1 > x_1$. Moreover, $d + d_3 + \cdots + d_s = d_{s+1} + \cdots + d_\ell$. Thus, by induction, the edges of the multigraph obtained from M by deleting e_1 and e_2 and adding an edge e between x_1 and y_2 can be decomposed into a collection of circuits. Adding back the edges e_1 and e_2 (which form a path between x_1 and y_2) in place of e, we obtain a suitable decomposition of the edges of M.

In the second case we may assume by relabelling that $y_1 = y_{s+1}$. If $x_1 = x_{s+1}$, then $d_1 = d_{s+1}$ and hence $d_2 + \cdots + d_s = d_{s+2} + \cdots + d_\ell$. By induction, the edges of M without e_1 and e_{s+1} can be decomposed into a collection of circuits. Adding the circuit consisting of e_1 and e_{s+1}, we obtain a decomposition of all edges of M. If $x_1 \neq x_{s+1}$, then without loss of generality $x_{s+1} > x_1$ and we have $d \coloneqq d_1 - d_{s+1} = x_{s+1} - x_1 \in D$ and $d + d_2 + \cdots + d_s = d_{s+2} + \cdots + d_\ell$. By induction, there is a decomposition of the edges of the multigraph obtained from M by deleting e_1 and e_{s+1} and adding a new edge e incident to x_1 and x_{s+1}. Adding back the edges e_1 and e_{s+1} (which form a path between x_1 and x_{s+1}) in place of e, we obtain a suitable decomposition of the edges of M.

For the 'moreover' part of the claim, note that if M has an odd number of edges, then the collection of circuits we obtain must contain a circuit with an odd number of edges. Any such circuit contains a cycle of odd length, which completes the proof of the claim. □

4.3.2 Proof of Theorem 4.1

Fix a graph F with at least one odd cycle, let $f = |V(F)|$ and let h be the odd girth of F. Let p be a prime large enough that $\left[\lfloor p/h \rfloor\right]$ has a B_h-subset S of size f that contains 1 and 2. Let the elements of S be $s_1 = 1, s_2 = 2, \ldots, s_f$ in ascending order and let the difference set of S be $D = \{s_j - s_i \colon i < j\}$. As F has at least one edge we may take a copy F^* of F with vertex set S which includes the edge $s_1 s_2$. We write

$$E = \{s_j - s_i \colon i < j \text{ and } s_i s_j \in E(F^*)\}$$

for the set of distances corresponding to edges of F^*. Since $s_1 = 1$ and $s_2 = 2$, we have $1 \in E$. Note that $E \subseteq D$. Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = E + p\mathbb{Z}$.

Let $\overrightarrow{X_n}$ be a digraph given by applying Lemma 4.7 with $g = h$. Recall that (a) the underlying graph X_n has chromatic number n, (b) the orientation $\overrightarrow{X_n}$ is acyclic, (c) for any pair of vertices u, v, there is at most one directed path from u to v in $\overrightarrow{X_n}$, and (d) for every cycle in X_n, the corresponding oriented cycle in $\overrightarrow{X_n}$ has at least h changes of direction. For vertices $u < v$ (in the reachability ordering), write $d(u, v)$ for the length of the unique directed path from u to v in $\overrightarrow{X_n}$. We now define a digraph $\vec{G}$ with vertex set $V(\overrightarrow{X_n})$ as follows.

$$\text{The edge } \overrightarrow{uv} \text{ is present if } u < v \text{ and } d(u, v) \in \overline{E}.$$

As before, the digraph $\vec{G}$ produced is acyclic. We now show that G, its underlying undirected graph, satisfies the properties stated in Theorem 4.1. Firstly since $1 \in E$, the graph X_n is a subgraph of G so $\chi(G) \geqslant n$. Next we check that every induced subgraph of G with sufficiently large chromatic number contains an induced copy of F. The argument is almost identical to the corresponding argument at the end of Section 4.2.2. We colour each edge uv of G (and $\vec{G}$), where $u < v$, with the colour $d(u, v) \bmod p$. As before, any monochromatic directed path in $\vec{G}$ of length pf gives rise to an induced copy of F. Let H be an induced subgraph of G which does not contain an induced copy of F. Each subgraph of H consisting of i-coloured edges is pf-colourable by Proposition 4.6 and so, by a product colouring, H is itself c'_F-colourable for $c'_F = (pf)^{|E|}$.

It remains to check that the odd girth of G is at least h. Suppose for a contradiction that G contains an odd cycle $v_0 v_1 \ldots v_{\ell-1}$ of length $\ell < h$. Since $\vec{G}$ contains no directed cycles, without loss of generality the path $v_{\ell-1} v_0 v_1$ in G is not a directed path in $\vec{G}$. We may further assume that since ℓ is odd, the path $v_0 v_1 v_2$ in G is a directed path in $\vec{G}$ from v_0 to v_2. For each edge $v_i v_{i+1}$ in this cycle (where here and

throughout we take subscript addition in the v_i to be modulo ℓ), there is a directed path $\overrightarrow{P_i}$ from v_i to v_{i+1} or from v_{i+1} to v_i in $\overrightarrow{X_n}$ depending on the direction of the edge between v_i and v_{i+1} in $\overrightarrow{G}$. For each i let P_i be the undirected path underlying $\overrightarrow{P_i}$ and concatenate these paths P_i in the natural way to obtain a walk W in X_n which begins at v_0, visits $v_1, \ldots, v_{\ell-1}$ in turn, and then finally returns to v_0.

For each traversal of an edge in the walk, when the direction of traversal is the same as the direction of the edge in $\overrightarrow{X_n}$, we consider the traversal of this edge in the walk to be 'coloured black'. When the two directions are different we consider the traversal of the edge to be 'coloured red'. For each of the paths P_i, the walk's traversals corresponding to P_i are all the same colour, and we have assumed that the colours for P_0 and P_1 are both black, so the walk changes colour at most $\ell - 2$ times (note that we do not include the change in colour between the end of the walk and the start).

Consider the subgraph L of G consisting of all vertices and edges contained in the walk W. We examine two cases based on whether or not L contains a cycle.

Case 1: Suppose that L is acyclic. Then L is a tree. Let $e = uv$ be an edge of L. The graph $L - e$ (i.e. the graph with vertex set $V(L)$ and edge set $E(L) \setminus \{e\}$) has exactly two components, one of which contains u and the other of which contains v. It is clear that the traversals of the edge e in the walk W alternate between traversing from u to v and traversing from v to u. Moreover, since the walk starts and ends at v_0, there are an equal number of traversals of e of in each direction. In other words, the total number of traversals of e coloured black is equal to the total number coloured red. Since this is true for all edges e, the total number of black traversals of edges in the walk is equal to the total number of red traversals.

Let $S_1 = \left\{0 \leqslant i \leqslant \ell - 1 \colon \overrightarrow{v_i v_{i+1}} \in E(\overrightarrow{G})\right\}$ and $S_2 = \left\{0 \leqslant i \leqslant \ell - 1 \colon \overrightarrow{v_{i+1} v_i} \in E(\overrightarrow{G})\right\}$ so that S_1 and S_2 form a partition of $\{0, 1, \ldots, \ell - 1\}$. For $0 \leqslant i \leqslant \ell - 1$, let d_i be the length of the path P_i. Then the total number of black traversals is equal to the sum of the d_i for $i \in S_1$ and the total number of red traversals is the sum of the d_i for $i \in S_2$. Hence by the above

$$\sum_{i \in S_1} d_i = \sum_{i \in S_2} d_i.$$

Since the edges $v_i v_{i+1}$ are all present in G, we have $d_i \in \overline{E}$ for all i. Let $\overline{d_i} \equiv d_i \bmod p$ be chosen so that $0 \leqslant \overline{d_i} \leqslant p - 1$. Then $\overline{d_i} \in E$. The sums $\sum_{i \in S_1} \overline{d_i}$ and $\sum_{i \in S_2} \overline{d_i}$ are

the same modulo p and each have fewer than h terms, all of which are at most p/h. It follows that, in fact,

$$\sum_{i \in S_1} \overline{d_i} = \sum_{i \in S_2} \overline{d_i}.$$

Noting that $E \subseteq D$, we can apply the cycle fact to show that the graph on vertex set S with edges between pairs of vertices at distances $\overline{d_0}, \ldots, \overline{d_{\ell-1}}$ contains an odd cycle of length at most ℓ. Since $\overline{d_0}, \ldots, \overline{d_{\ell-1}} \in E$, by the definition of E and the fact that S is a Sidon set, F^* contains a cycle of this length, which is a contradiction. This completes the analysis of the case where L does not contain a cycle.

Case 2: Suppose that L contains a cycle. Let Γ be the first cycle produced by the walk, and label its vertices as $c_0, c_1, \ldots, c_{r-1}$ in cyclic order around Γ, where c_0 is the vertex amongst these that W arrives at first. Let $C = \{c_0, \ldots, c_{r-1}\}$. We will show that Γ contains fewer than h changes of direction, which contradicts property (d) of $\overrightarrow{X_n}$. In what follows, we are only concerned with the portion of walk W from its first visit to c_0 to the point when the cycle Γ is formed. Let W' be this segment of W. In particular, c_0 is the first vertex of W' and the final edge traversal in W' completes the cycle Γ. Let this final edge traversal in W' be from c_a to c_b, where $b \equiv a \pm 1 \bmod r$.

Consider the subgraph of L traversed by W' up to the point just before it traverses its final edge $c_a c_b$. This graph is acyclic and contains every edge of Γ except $c_a c_b$. In other words, it is the union of the path $\Gamma - c_a c_b$ and r disjoint trees rooted at $c_0, \ldots, c_{r-1}$ respectively. From this, it is straightforward to see that there exist vertices $\gamma_0, \ldots, \gamma_{s-1}$ in C, with $\gamma_0 = c_0$ and γ_{j+1} a neighbour of γ_j in Γ, such that W' has the following form. It is a (possibly empty) walk from γ_0 to itself avoiding $C \setminus \{\gamma_0\}$, followed by a traversal of the edge $\gamma_0\gamma_1$, followed by a (possibly empty) walk from γ_1 to itself avoiding $C \setminus \{\gamma_1\}$, followed by a traversal of the edge $\gamma_1\gamma_2$, and so on, concluding with a traversal of the edge $\gamma_{s-2}\gamma_{s-1}$, which completes the cycle Γ.

The vertices $\gamma_0, \gamma_1, \ldots, \gamma_{s-1}$ form a walk W'' on Γ traversing all the edges of Γ. Let the traversals of each edge have the same colour in W'' as the corresponding traversals in W'. Note that W'' changes colour at most as many times as W'. We say that a vertex $c_i \in C$ is *crossed* by the walk W'' if c_{i-1}, c_i, c_{i+1} are consecutive vertices of W'' in either order. If there are two vertices of C not crossed by W'', then W'' does not traverse all edges of Γ, so there is at most one vertex not crossed by W''.

If the edges $c_{i-1}c_i$ and $c_i c_{i+1}$ are a change of direction, then when W'' crosses c_i the traversals of these two edges are of different colours. It follows that every time W'' crosses c_i, it changes colour between the traversals of the two edges. Hence,

since W'' changes colour at most $h-2$ times and all but at most one of the vertices in C are crossed, there are at most $h-1$ changes of direction in Γ. We have thus obtained a cycle in $\overrightarrow{X_n}$ with fewer than h changes of direction, which gives the required contradiction. Therefore, the odd girth of G is at least h.

4.4 Tournaments

We now prove Theorem 4.4, which is restated below. Our strategy is akin to that used by Alon, Pach, and Solymosi in [4]. Recall that a *tournament* is a complete graph in which each edge has an orientation, and that a tournament is *transitive* if, for all distinct vertices u, v, w, the presence of the edges $\overrightarrow{uv}$ and $\overrightarrow{vw}$ in the tournament implies the presence of the edge $\overrightarrow{uw}$. The *chromatic number* $\chi(T)$ of a tournament T is the smallest possible number of parts in a partition of its vertex set in which each part induces a transitive tournament. We will prove the following corollary of Theorem 1.8.

Theorem 4.4. *For every tournament T, there is a constant, C_T, for which there exist tournaments S of arbitrarily large chromatic number such that every T-free subtournament of S has chromatic number at most C_T.*

Note that there are certain tournaments which are contained in every tournament of sufficiently large chromatic number; all such tournaments were described explicitly in [20]. Before proving the corollary, we state two results we will need. The first is the following standard corollary of Dilworth's theorem [42].

Lemma 4.13. *Let $(a_1, a_2, \ldots, a_k)$ be a sequence of distinct elements and let $<$ be a strict total order on them. If this sequence has no decreasing (with respect to $<$) subsequence of length greater than m, then it can be partitioned into at most m increasing subsequences.*

We will also use the following result originally proved by Rödl and Winkler in [123]. An *ordering* of a graph or digraph G is a strict total ordering of $V(G)$, and an *ordered (di)graph* is a pair $(G, <)$ where G is a (di)graph and $<$ is an ordering of G.

Theorem 4.14 ([123])**.** *For any ordered graph $(B, <)$, there exists a graph F such that, for every ordering $<'$ of F, $(B, <)$ is an ordered induced subgraph of $(F, <')$.*

We are now ready to prove Theorem 4.4.

Proof of Theorem 4.4. Fix an ordered tournament $(T, <)$ with vertices $x_1 < \cdots < x_{|T|}$. Let $(B, <)$ be the ordered back-edge graph of $(T, <)$, that is, the ordered graph on the same vertex set as T and with the same ordering as T, where $x_i x_j \in E(B)$ for $i < j$ exactly when $\overrightarrow{x_j x_i} \in E(T)$. Let F be the graph obtained by applying Theorem 4.14 to $(B, <)$. Now apply Theorem 1.8 to F. We obtain a constant $C_T = c_F$ and a graph L of (arbitrarily large) chromatic number n with $\omega(L) = \omega(F) = \omega$ such that every induced subgraph of L with chromatic number greater than C_T contains an induced copy of F. Fix an arbitrary ordering $<'$ of L and define an ordered tournament $(S, <')$ with the same vertex set as L and the same ordering as L by orienting $\overrightarrow{xy}$ for $x <' y$ if $xy \notin E(L)$ and $\overrightarrow{yx}$ otherwise (so that $(L, <')$ is the ordered back-edge graph of $(S, <')$).

We now show that $\chi(S) \geqslant n/\omega$, and hence that $\chi(S)$ is arbitrarily large. Consider a transitive subtournament A in S and write its vertices as a sequence $(a_1, a_2, \ldots, a_{|A|})$ such that $\overrightarrow{a_i a_j} \in E(S)$ for each $i < j$. Observe that this sequence has no decreasing (with respect to $<'$) subsequence of length $\omega + 1$ as this would give rise to an $(\omega + 1)$-clique in L. Hence, by Lemma 4.13, we may partition the sequence into at most ω increasing subsequences. Each of these corresponds to an independent set in L and, hence, $\omega \cdot \chi(S) \geqslant n$ as required.

It remains to show that every subtournament R of S with $\chi(R) > C_T$ contains a copy of T. If R is such a subtournament, then clearly the induced subgraph R' of L with the same vertex set as R has $\chi(R') > C_T$. Hence, by our assumptions on L, R' contains an induced copy of F. Consider the ordering of F given by the restriction of $<'$ to this copy. By the construction of F, under this ordering it contains an induced ordered copy of $(B, <)$ which guarantees the existence of a copy of T in R, as required. This completes the proof of the theorem. □

4.5 Hypergraphs

In this section we prove Theorem 4.5, which is restated below. Recall that in this chapter all hypergraph edges have size at least two. For a hypergraph $\mathcal{G}$ and a set $X \subseteq V(\mathcal{G})$ of vertices of $\mathcal{G}$, the *subhypergraph of* $\mathcal{G}$ *induced on* X, denoted $\mathcal{G}[X]$, is the hypergraph with vertex set X and edge set $\{e \colon e \subseteq X, e \in E(\mathcal{G})\}$. For hypergraphs $\mathcal{F}$ and $\mathcal{G}$, we will say that $\mathcal{G}$ is *$\mathcal{F}$-free* if it does not contain $\mathcal{F}$ as an induced subhypergraph. The *chromatic number* of a hypergraph is the minimum number of colours required to colour its vertices so that no edge is monochromatic, and a hypergraph

is said to be *strongly t-colourable* if its vertices can be t-coloured such that no edge contains two vertices of the same colour. We say that a hypergraph *covers* a pair of vertices if it has an edge containing both of them.

Theorem 4.5. *For every hypergraph $\mathcal{F}$ with at least one edge, there is a constant, $c_{\mathcal{F}}$, for which there exist hypergraphs $\mathcal{G}$ of arbitrarily large chromatic number such that every $\mathcal{F}$-free induced subhypergraph of $\mathcal{G}$ is strongly $c_{\mathcal{F}}$-colourable. Moreover, we can take $\mathcal{G}$ such that if $\mathcal{G}$ covers every pair from some $X \subseteq V(\mathcal{G})$, then $\mathcal{G}[X]$ is an induced subhypergraph of $\mathcal{F}$.*

The proof of Theorem 4.5 resembles those of Theorems 1.8 and 4.1 in that we start with an oriented base (hyper)graph with large chromatic number and certain 'reachability' properties, to which we add edges based on 'distances' between vertices modulo some prime p. As before, we will ascertain which edges to add to the base hypergraph by placing a copy of $\mathcal{F}$ on a B_3-set, an approach reminiscent of that used in the earlier proofs. Once we have constructed $\mathcal{G}$, the arguments we use to show it has the desired properties also bear strong similarities to those used above.

In our construction it will be helpful to use the following standard notion of the girth of a hypergraph. For our purposes, a *cycle of length $\ell \geqslant 2$* in a hypergraph is a sequence of distinct vertices $v_0, \dots, v_{\ell-1}$ and a sequence of distinct edges $e_0, \dots, e_{\ell-1}$ such that $v_i \in e_i \cap e_{i+1}$ for all i, where the indices are taken modulo ℓ. Note that in particular we consider two distinct edges with at least two vertices in their intersection to give rise to a cycle of length 2. The *girth* of a hypergraph is the length of its shortest cycle.

We now introduce an analogue of reachability and distances between vertices in the hypergraph setting. Given a hypergraph $\mathcal{G}$ and a total order $\prec$ on its vertex set, we construct the *$\prec$-digraph* $\vec{G}$ on vertex set $V(\mathcal{G})$ by adding to $E(\vec{G})$ the directed edges $\overrightarrow{v_1v_2}, \overrightarrow{v_2v_3}, \dots, \overrightarrow{v_{a-1}v_a}$ for each edge $e = \{v_1 \prec v_2 \prec \cdots \prec v_a\}$ of $\mathcal{G}$. The *$\prec$-graph* G is the underlying undirected graph of $\vec{G}$. If $\mathcal{G}$ has girth at least 3, then no two distinct edges intersect in more than one vertex, and it follows that every pair of distinct vertices has at most one directed edge between them in $\vec{G}$. Also, in this setting every edge in $\vec{G}$ comes from a unique edge of $\mathcal{G}$.

The base hypergraph for our construction should have large chromatic number as well as an ordering $\prec$ on its vertex set such that its $\prec$-digraph is suitable for defining reachability and distance. The following lemma states that such hypergraphs exist. Similarly to Lemma 4.7, the proof given below is based on an observation of

Nešetřil [110] and replaces a more complicated argument analysing a construction of his and Rödl's [114]. For $k \geqslant 2$ a hypergraph is *k-uniform* if every edge contains exactly k vertices.

Lemma 4.15. *For all integers $k \geqslant 2$, $g \geqslant 3$, and $n \geqslant 2$, there exists a k-uniform hypergraph $\mathcal{Y}_n$ with chromatic number n and girth at least g, and an order $\prec$ on its vertex set such that, for every cycle in its $\prec$-graph Y_n, the corresponding oriented cycle in its $\prec$-digraph $\overrightarrow{Y_n}$ has at least g changes of direction.*

Proof. Nešetřil and Rödl [114] proved that for all integers $k \geqslant 2$, $g \geqslant 3$, and $n \geqslant 2$ there is a k-uniform hypergraph with chromatic number n and girth at least g which is strongly a-colourable[1] for $a = (k-1)n + 1$.

Fix integers $k \geqslant 2$, $g \geqslant 3$, and $n \geqslant 2$, and let $a = (k-1)n + 1$. By the result of Nešetřil and Rödl, there is a k-uniform hypergraph $\mathcal{Y}_n$ with chromatic number n and girth greater than $(g-1)(a-1)$ which is strongly a-colourable. Fix a strong a-colouring $c\colon V(\mathcal{Y}_n) \to \{1, \dots, a\}$ of $\mathcal{Y}_n$. Let $\prec$ be any ordering of $V(\mathcal{Y}_n)$ in which every vertex of colour i precedes every vertex of colour j whenever $i < j$. Let Y_n and $\overrightarrow{Y_n}$ respectively be the $\prec$-graph and $\prec$-digraph of $\mathcal{Y}_n$.

Consider any edge $\overrightarrow{uv}$ in $\overrightarrow{Y_n}$. There is some edge of $\mathcal{Y}_n$ containing both u and v and so $c(u) \neq c(v)$. By the choice of $\prec$, we must have $c(u) < c(v)$. Letting $u_1 u_2 \dots u_t$ be any directed path in $\overrightarrow{Y_n}$, by the preceding discussion we have $1 \leqslant c(u_1) < \dots < c(u_t) \leqslant a$. Hence, the length of any directed path in $\overrightarrow{Y_n}$ is at most $a - 1$.

Consider any cycle C in Y_n: each edge of this cycle comes from an edge of a cycle in $\mathcal{Y}_n$ and so C has length greater than $(g-1)(a-1)$. Every directed path in $\overrightarrow{Y_n}$ has length at most $a - 1$, so C has at least g changes of direction in $\overrightarrow{Y_n}$, as required. □

We are now ready to prove Theorem 4.5.

Proof of Theorem 4.5. Fix a hypergraph $\mathcal{F}$ with at least one edge. Let f be the number of vertices of $\mathcal{F}$ and let $n \geqslant 2$. Let $m \geqslant 2$ be the minimum edge size of $\mathcal{F}$. Let $\mathcal{Y}_n$ be an m-uniform hypergraph of chromatic number at least n and girth at least four given by Lemma 4.15, with an ordering $\prec$ on its vertex set such that in the associated digraph $\overrightarrow{Y_n}$ every cycle has at least three changes of direction. In particular, in $\overrightarrow{Y_n}$ there are no directed cycles and there is at most one directed path

[1] Although the strongly a-colourable condition does not appear in their theorem statement, it is emphasised in their proof.

between any pair of vertices. If there is a directed path from u to v in $\overrightarrow{Y_n}$, then we write $u < v$ and define the distance between u and v, written $d(u, v)$, to be the length of this path.

We will define a hypergraph $\mathcal{G}$ satisfying the conditions of Theorem 4.5 by adding edges to $\mathcal{Y}_n$ as follows. Start by picking a prime p such that $[\lfloor p/(4m) \rfloor]$ contains a B_3-set S' of size f, then define $S = 2m \cdot S' = \{2ms' \colon s' \in S'\} \subseteq [(p-1)/2]$. Let the elements of S be $s_1, \ldots, s_f$ in ascending order and let $D = \{s_j - s_i \colon i < j\}$ be the difference set of S. Note that all elements of D are positive multiples of $2m$ and are at most $(p-1)/2$.

Taking a copy $\mathcal{F}^*$ of $\mathcal{F}$ with vertex set S, we define

$$E = \{(s_{i_2} - s_{i_1}, s_{i_3} - s_{i_2}, \ldots, s_{i_a} - s_{i_{a-1}}) : i_1 < \cdots < i_a \text{ and } \{s_{i_1}, s_{i_2}, \ldots, s_{i_a}\} \in E(\mathcal{F}^*)\}$$

to be the analogue of the sets of allowable distances from the proofs of Theorems 1.8 and 4.1. Note that every entry of every tuple in E is in D. Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = \{(d_1 + k_1 p, \ldots, d_{a-1} + k_{a-1} p) \colon k_1, \ldots, k_{a-1} \in \mathbb{Z} \text{ and } (d_1, \ldots, d_{a-1}) \in E\}$. Construct a hypergraph $\mathcal{G}$ by *adding* edges to $\mathcal{Y}_n$ as follows.

$$\text{The edge } \{v_1, v_2, \ldots, v_a\} \text{ is added if } v_1 < v_2 < \cdots < v_a \text{ and}$$
$$(d(v_1, v_2), d(v_2, v_3), \ldots, d(v_{a-1}, v_a)) \in \overline{E}.$$

There are two types of edges in $\mathcal{G}$: those which appear in $\mathcal{Y}_n$ and the newly added ones. By construction, if two vertices are in a common edge, then they are $<$-comparable. The edges which appear in $\mathcal{Y}_n$ are of the form $\{v_1, \ldots, v_m\}$ where $d(v_i, v_{i+1}) = 1$ for all i. Hence, if vertices $u < v$ are both in some edge of $\mathcal{Y}_n$, then $d(u, v) \in [m-1]$. On the other hand, if $u < v$ are both in an edge that has been newly added, then $d(u, v) \in \overline{D}$ and so $d(u, v) \geqslant 2m$. In particular, if vertices $u < v$ are both in some edge of $\mathcal{G}$, then $d(u, v)$ is in $\overline{D} \cup [m-1]$ and so is not divisible by p.

We will now show that this hypergraph satisfies the conditions of Theorem 4.5. Firstly, since it contains $\mathcal{Y}_n$, it certainly has chromatic number at least n. Next, we will show that every $\mathcal{F}$-free induced subhypergraph of $\mathcal{G}$ has bounded strong chromatic number. We begin by defining L to be the graph with vertex set $V(\mathcal{G})$ where uv is an edge if u and v are in some common edge of $\mathcal{G}$. For each edge uv of L, u and v are $<$-comparable. Orient the edge from u to v if $u < v$, and from v to u otherwise. Denote the resulting digraph by $\overrightarrow{L}$.

All edges $\overrightarrow{uv}$ of $\overrightarrow{L}$ have endpoints which satisfy $d(u,v) \in \overline{D} \cup [m-1]$, and we colour the edges of $\overrightarrow{L}$ (and L) by the residue modulo p of this distance. We claim that the vertex set of a monochromatic directed path of length pf in $\overrightarrow{L}$ has a subset which induces a copy of $\mathcal{F}$ in $\mathcal{G}$. Indeed, by repeatedly jumping at most $p-1$ steps along such a path we can find a sequence of vertices $u_1 < u_2 < \cdots < u_f$ in the path such that $d(u_i, u_{i+1}) \equiv s_{i+1} - s_i \bmod p$ for all i. Note that, for $i < j$, $d(u_i, u_j) \equiv s_j - s_i \bmod p$ and so $d(u_i, u_j) \in \overline{D}$. Hence, $d(u_i, u_j)$ is never 1 and so $\{u_1, \ldots, u_f\}$ does not contain an edge of $\mathcal{Y}_n$. Thus, by the construction of $\mathcal{G}$, the vertices $u_1, \ldots, u_f$ induce a copy of $\mathcal{F}$, as claimed.

Now let $\mathcal{H}$ be an $\mathcal{F}$-free induced subhypergraph of $\mathcal{G}$. By the above, for each $i \in D \cup [m-1]$ there is no i-coloured directed path of length pf in $\overrightarrow{L}[V(\mathcal{H})]$, so the i-coloured subgraph of $L[V(\mathcal{H})]$ is (pf)-colourable by Proposition 4.6. A product colouring demonstrates that $L[V(\mathcal{H})]$ is $(pf)^{p-1}$-colourable. This colouring is a strong colouring of $\mathcal{H}$. Hence we may take $c_{\mathcal{F}} = (pf)^{p-1}$, which is a constant that only depends on $\mathcal{F}$.

It remains to show the moreover part of the statement. Let $X = \{v_1, \ldots, v_t\}$ be a set of vertices every pair of which is covered by $\mathcal{G}$. If $t \leqslant 1$, then the result is clear so assume $t \geqslant 2$. It follows from the fact that $\mathcal{G}$ covers every pair from X that each pair is $<$-comparable; we may assume that $v_1 < \cdots < v_t$. Define $d_i = d(v_i, v_{i+1})$ for all i and note that since $d(v_i, v_j) \in \overline{D} \cup [m-1]$ for all $i < j$, all sums of the form

$$\sum_{i_1 \leqslant j \leqslant i_2} d_j \tag{4.4}$$

with $1 \leqslant i_1 \leqslant i_2 \leqslant t-1$ are in $\overline{D} \cup [m-1]$.

Suppose that $d_i \leqslant m-1$ for some i. If $i < t-1$, then d_{i+1} and $d_i + d_{i+1}$ differ by at most $m-1$. This cannot occur if $d_{i+1} \in \overline{D}$, so we conclude that $d_{i+1} \leqslant m-1$. Similarly, if $i > 1$ then $d_{i-1} \leqslant m-1$. Repeating this argument we find that $d_1, \ldots, d_{t-1} \in [m-1]$. Since the minimal positive element of $\overline{D}$ is at least $2m$, it follows that $d_1 + d_2 \notin \overline{D}$, so $d_1 + d_2 \in [m-1]$. Applying this repeatedly we find that all sums of the form in (4.4) are in $[m-1]$, and hence no pair of vertices from X appear together in an edge of $\mathcal{G}$ which is not in $\mathcal{Y}_n$. In particular, $\mathcal{Y}_n$ covers all pairs from X.

Let e be an edge of $\mathcal{Y}_n$ containing both v_1 and v_2, and suppose there is some v_a not in e. Let e_1, e_2 be edges of $\mathcal{Y}_n$ such that e_i contains v_a and v_i. If $e_1 = e_2$, then the edges e and e_1 form a 2-cycle in $\mathcal{Y}_n$, while if $e_1 \neq e_2$, then the edges e, e_1, e_2 form a 3-cycle in $\mathcal{Y}_n$. As $\mathcal{Y}_n$ has girth at least four, neither of these cases is possible. We

deduce that $X \subseteq e$, and in particular that $t \leqslant m$. If $t < m$, then the subhypergraph of $\mathcal{G}$ induced on X has no edges and thus is clearly an induced subhypergraph of $\mathcal{F}$. If $t = m$, then the subhypergraph of $\mathcal{G}$ induced on these vertices has a single edge of size m, which again is clearly an induced subhypergraph of $\mathcal{F}$.

On the other hand, if $d_i > m - 1$ for all i, then all the sums of the form in (4.4) are in $\overline{D}$. Writing $\overline{d_i}$ for the residue modulo p of d_i, we have $\overline{d_i} \in D$ for all i and all sums of the form

$$\sum_{i_1 \leqslant j \leqslant i_2} \overline{d_j}$$

with $1 \leqslant i_1 \leqslant i_2 \leqslant t - 1$ are in $\overline{D}$. We can deduce from the fact that $D \subseteq [(p-1)/2]$ that in fact all of these sums are in D. It now follows from the clique fact that there exist $b_1 < \cdots < b_t$ in S such that $b_{i+1} - b_i = \overline{d_i}$ for all i. Since $\overline{d_i} \neq 1$ for all i, the subhypergraph of $\mathcal{G}$ induced on X does not include any edge of $\mathcal{Y}_n$, so by the construction of $\mathcal{G}$ and the fact that S is a Sidon set, the subhypergraph of $\mathcal{G}$ induced on X is isomorphic to that of $\mathcal{F}^*$ induced on $\{b_1, \ldots, b_t\}$, and the claim follows. □

4.6 Extensions to infinite families of graphs

Recall that the *disjoint union* of a family of graphs is the (possibly infinite) graph consisting of pairwise vertex-disjoint copies of the graphs in the family with no edges between copies. As noted at the start of this chapter, for a finite family of graphs $\mathcal{F}$ (at least one of which contains an edge), an easy consequence of Theorem 1.8 is that there is a constant $c_{\mathcal{F}}$ for which there exist graphs G of arbitrarily large chromatic number and the same clique number as the largest clique number of a graph in $\mathcal{F}$ such that every induced subgraph of G with chromatic number greater than $c_{\mathcal{F}}$ contains every member of $\mathcal{F}$ as an induced subgraph. Taking a disjoint union of the graphs G of arbitrarily large chromatic number gives the following corollary.

Corollary 4.16. *Let $\omega \geqslant 2$, and suppose that $\mathcal{F}$ is a finite family of finite graphs all with clique number at most ω. There exists a constant $c_{\mathcal{F}}$ and an infinite graph G of infinite chromatic number and clique number at most ω such that every induced subgraph of G with chromatic number greater than $c_{\mathcal{F}}$ contains every member of $\mathcal{F}$ as an induced subgraph.*

Does this phenomenon occur for infinite families of graphs? To this end we will say that a (possibly infinite) family of graphs $\mathcal{F}$ is *durable* if there exists an infinite graph G with infinite chromatic number such that every infinite-chromatic induced

subgraph of G contains every member of $\mathcal{F}$ as an induced subgraph. If such a G exists we will call it a *witness* of $\mathcal{F}$'s durability. We note a few properties of witnesses and durability.

- A subfamily of a durable family is durable.
- Every infinite-chromatic induced subgraph of a witness is also a witness.
- Every witness contains every graph in $\mathcal{F}$ as an induced subgraph.

Corollary 4.16 says that every finite family of graphs is durable. Moreover, if at least one graph in the finite family contains an edge, then there is a witness whose clique number is no bigger than the largest clique number of a member of the family.

We are interested in which infinite families of finite graphs are durable. There certainly are some durable infinite families, for example any family whose members are all disjoint unions of complete graphs is durable as witnessed by the disjoint union of K_1, K_2, K_3, ... This example suggests that for a durable family with unbounded chromatic number, the disjoint union of the members of the family might be a witness of the family's durability. If this were the case, then this disjoint union would be a minimal witness. The following result confirms this hypothesis.

Theorem 4.17. *Let $\mathcal{F}$ be a countable family of finite graphs. Let $G_{\mathcal{F}}$ be the disjoint union of the members of $\mathcal{F}$.*

(a) *If $\mathcal{F}$ is durable, then every witness of $\mathcal{F}$'s durability contains $G_{\mathcal{F}}$ or K_∞ as an induced subgraph.*
(b) *If $\{\chi(F)\colon F \in \mathcal{F}\}$ is unbounded, then $\mathcal{F}$ is durable if and only if $G_{\mathcal{F}}$ is a witness of $\mathcal{F}$'s durability.*

Theorem 4.17(b) shows that if a family of unbounded chromatic number is durable, then it has a witness whose clique number is no bigger than the largest clique number of a member of the family. An immediate application of Theorem 4.17 is Theorem 4.2: the family, $\mathcal{F}_g$, of finite graphs of girth at least g is not durable. Indeed, this family has unbounded chromatic number and so if it were durable, then $G_{\mathcal{F}_g}$ would witness this. However, the disjoint union of graphs of girth at least $g+1$ is an infinite-chromatic induced subgraph of $G_{\mathcal{F}_g}$ that does not contain the cycle of length g.

The following lemma will be needed in the proof of Theorem 4.17.

Lemma 4.18. *Every infinite graph either contains an infinite clique, has bounded clique number or contains the disjoint union of all finite cliques as an induced subgraph.*

Proof. Let G be an infinite graph with unbounded clique number. Let $(a_i)_{i\geqslant 1}$ be a sequence of natural numbers in which each term is taken to be large relative to all previous terms. Fix a copy of K_{a_1} in G on vertex set V_1. By identifying a copy of $K_{a_1+a_2}$ in G, we can find a copy of K_{a_2} on vertex set V_2 disjoint from V_1. Continuing in this manner, we obtain a sequence $(V_i)_{i\geqslant 1}$ of pairwise disjoint sets of vertices of G such that V_i induces a copy of K_{a_i} for each i.

Select a vertex $v \in V_1$, and for each $i \geqslant 2$ remove at most half of the vertices from V_i so that v is either adjacent to every remaining vertex or not adjacent to any of them. Repeat this process for each of the other vertices in V_1, removing at most half of the vertices currently in each V_i at each step, so that at the end of the process the remaining sets of vertices $V_i' \subseteq V_i$ satisfy $|V_i'| \geqslant |V_i|/2^{a_1}$ and have the property that for all $i \geqslant 2$ and $v \in V_1$, v is either adjacent to every vertex in V_i' or not adjacent to any of them.

For each $i \geqslant 2$ we can now define a vector $b_i \in \{0,1\}^{a_1}$ by letting the jth entry be 1 if the jth vertex of V_1 is adjacent to every vertex in V_i' and 0 otherwise. There are only finitely many such vectors, so there exists a sequence (i_k) such that all vectors b_{i_k} agree. By relabelling, we may therefore assume that all vectors b_i are equal to some b. Observe that since each a_i is arbitrarily large relative to $a_1, \dots, a_{i-1}$, after this relabelling we may still assume that each set V_i' is arbitrarily large relative to $V_2', \dots, V_{i-1}'$ and to a_1. If at least half the entries in b are 1's, then by considering the vertices in V_1 corresponding to 1 entries, for some $t \geqslant a_1/2$ we obtain a copy of K_t in G, every vertex of which is adjacent to every vertex in $\bigcup_{i\geqslant 2} V_i'$. Otherwise, for some $t \geqslant a_1/2$ we obtain a copy of K_t none of whose vertices have a neighbour in this union.

Repeat this process for V_2', then what remains of V_3', and so on. If the a_i are chosen appropriately, then it follows that there exists a sequence of pairwise disjoint sets of vertices of G, $(W_n)_{n\geqslant 1}$, such that W_n induces a copy of K_n in G and the possible edges between W_n and $\cup_{i>n} W_i$ are either all present in G or all not present in G. If these edges are present for infinitely many n, then G contains an infinite clique. If the edges are present for only finitely many n, then there is an induced subgraph of G consisting of a disjoint union of arbitrarily large cliques, which in particular contains the disjoint union of all finite cliques as an induced subgraph. □

We are ready to prove Theorem 4.17. For a vertex v in a (possibly infinite) graph G, let $\Gamma(v) \coloneqq \{u \in V(G) \colon uv \in E(G)\}$ be the *neighbourhood* of v in G. Similarly, for

$V \subseteq V(G)$ we write $\Gamma(V) \coloneqq \{u \in V(G) \colon uv \in E(G) \text{ for some } v \in V\}$.

Proof of Theorem 4.17. We first show that (a) implies (b) in the statement of the theorem. Suppose that $\mathcal{F}$ is durable with $\{\chi(F) \colon F \in \mathcal{F}\}$ unbounded, and let G be a witness of $\mathcal{F}$'s durability. By part (a), G contains $G_{\mathcal{F}}$ or K_∞ as an induced subgraph. In the former case, since $\{\chi(F) \colon F \in \mathcal{F}\}$ is unbounded, $G_{\mathcal{F}}$ has infinite chromatic number and is therefore a witness by one of the properties of witnesses noted above. In the latter case, again by the properties noted above, K_∞ is a witness of $\mathcal{F}$'s durability and hence every member of $\mathcal{F}$ is an induced subgraph of K_∞. Thus, $\mathcal{F}$ is a family of cliques and it follows from the fact that $\{\chi(F) \colon F \in \mathcal{F}\}$ is unbounded that $G_{\mathcal{F}}$ is a witness.

It remains to prove the first part of the theorem. Suppose that $\mathcal{F}$ is durable and let G be a witness. By Lemma 4.18, G either contains an infinite clique, has bounded clique number or contains the disjoint union of all finite cliques as an induced subgraph. In the first case we are done. In the third case, the disjoint union of all finite cliques is a witness of $\mathcal{F}$'s durability, so all members of $\mathcal{F}$ are disjoint unions of cliques. In particular, the disjoint union of all finite cliques contains $G_{\mathcal{F}}$ as an induced subgraph and so we are done.

Hence, we are left with the case where G has bounded clique number. We may assume that the graph induced on the neighbourhood of each vertex of G is finitely colourable. Indeed, otherwise there exists $v \in V(G)$ such that $G[\Gamma(v)]$ is an infinite-chromatic induced subgraph of G. In this case $G[\Gamma(v)]$ is a witness of $\mathcal{F}$'s durability which has clique number strictly less than that of G. Repeating this process finitely many times, we obtain an infinite-chromatic induced subgraph of G that is a witness of $\mathcal{F}$'s durability and in which the neighbourhood of every vertex is finitely colourable.

Enumerate $\mathcal{F}$ as $(F_i)_{i \geqslant 1}$. Since G is a witness, G contains an induced copy of F_1, say on vertex set V_1. By assumption, the chromatic number of the subgraph of G induced on $V_1 \cup \Gamma(V_1)$ is finite. It follows that $G - (V_1 \cup \Gamma(V_1))$ is an infinite-chromatic induced subgraph of G and therefore contains every graph in $\mathcal{F}$ as an induced subgraph. We may now find an induced copy of F_2 in this graph, say on vertex set V_2, and remove every vertex in V_2 or its neighbourhood from $G - (V_1 \cup \Gamma(V_1))$ to obtain a new infinite-chromatic induced subgraph of G. Continuing in this manner we obtain an induced subgraph of G consisting of the disjoint union of the members of $\mathcal{F}$, that is, G contains $G_{\mathcal{F}}$ as an induced subgraph. □

Theorem 4.17 gives an effective test for whether a family of unbounded chromatic number is durable. Still, it is not clear whether there are durable families of unbounded chromatic number whose members are not all disjoint unions of cliques. Theorem 4.17 shows that the existence of such a family is equivalent to a positive answer to the following, as explained below.

Question 4.19. *Do there exist sequences of connected finite graphs $(F_i)_{i\geqslant 1}$, not all of which are cliques, and of positive integers $(c_i)_{i\geqslant 1}$ such that $\chi(F_i) \to \infty$ and every F_i-free induced subgraph of F_j is c_i-colourable for all $i < j$?*

Indeed, given sequences $(F_i)_{i\geqslant 1}$ and $(c_i)_{i\geqslant 1}$ as in Question 4.19, the disjoint union of the family $(F_i)_{i\geqslant 1}$ has infinite chromatic number and the property that all of its F_i-free induced subgraphs are $\max\{c_i, \chi(F_1), \dots, \chi(F_i)\}$-colourable. It follows that $(F_i)_{i\geqslant 1}$ is a durable family of graphs with unbounded chromatic number, not all of whose members are disjoint unions of cliques.

On the other hand, if $\mathcal{F}$ is a durable family of unbounded chromatic number whose members are not all disjoint unions of cliques, then let F_1 be a non-clique connected component of a graph in $\mathcal{F}$ and let $(F_i)_{i\geqslant 2}$ be an arbitrary sequence of connected components of graphs in $\mathcal{F}$ with the property that $\chi(F_i) \to \infty$. Note that $(F_i)_{i\geqslant 1}$ is a durable family: any witness of the durability of $\mathcal{F}$ is also a witness of this family's durability. By Theorem 4.17, the disjoint union of the F_i is a witness of their durability. It follows that for each i there is a constant c_i such that every F_i-free induced subgraph of F_j is c_i-colourable for all $i < j$.

We now turn to infinite families with bounded chromatic number. The next result provides some durable families of forests. A graph F is *χ-bounding* if the class of F-free graphs is χ-bounded. It is folklore that any χ-bounding graph must be a forest and the celebrated Gyárfás–Sumner conjecture [65, 139] asserts the converse. We remark that a beautiful proof of Gyárfás [66] shows that every path is χ-bounding and so the following theorem, combined with the first bullet point above, implies Theorem 4.3.

Theorem 4.20. *The family of χ-bounding graphs is durable.*

Proof. Let G be a triangle-free graph with infinite chromatic number. Let H be an induced subgraph of G with infinite chromatic number. Fix a χ-bounding graph F and note that there is some constant c_F such that every triangle-free F-free graph is c_F-colourable. But H is triangle-free and has chromatic number greater than c_F,

so it contains F as an induced subgraph. Hence, G witnesses the durability of the family of χ-bounding graphs. □

It is possible that every family with bounded chromatic number is durable (although we suspect this is not the case). An answer to the following very natural question would resolve this.

Question 4.21. *Is the family of k-colourable graphs durable? That is, for each positive integer k, is there a graph G_k with infinite chromatic number such that every infinite-chromatic induced subgraph of G_k contains every finite k-colourable graph as an induced subgraph?*

It would be natural to ask in addition for the graph G_k to have clique number k to align with Corollary 4.16.

4.7 Conclusion

Having considered tournaments, hypergraphs, and infinite families of graphs, we now focus on the setting of graphs. As discussed in Section 1.3, Theorem 1.8 shows that the class $\mathcal{C}_r$ of K_r-free graphs has a far stronger property than just being a vertex Ramsey class, and it is natural to ask which other families have this property.

Question 4.22. *Which hereditary graph classes $\mathcal{C}$ have the property that for every $F \in \mathcal{C}$ there is a constant, c_F, for which there exist graphs $G \in \mathcal{C}$ of arbitrarily large chromatic number such that every F-free induced subgraph of G is c_F-colourable?*

A good starting point for tackling Question 4.22 would be to resolve the case when $\mathcal{C}$ is determined by a single excluded graph.

A case of particular interest is when $\mathcal{C}$ is the class of graphs with girth at least g for some $g \geqslant 3$. Theorem 4.1 asserts that the class of graphs with *odd girth* at least g has the stated property. As mentioned previously, we conjecture that the same is true for girth.

Conjecture 4.23. *For every graph F with at least one cycle, there is a constant, b_F, for which there exist graphs G of arbitrarily large chromatic number and the same girth as F such that every F-free induced subgraph of G is b_F-colourable.*

For this conjecture to hold, the class of graphs with girth at least g must certainly be vertex Ramsey. Happily, since cycles are 2-connected this follows from a theorem of Nešetřil and Rödl [112].

The graphs $\overrightarrow{X_n}$ and X_n described in Section 4.1 can be chosen to not contain short cycles. However, if we continue to build our graphs by adding edges $\overrightarrow{uv}$ for $u < v$ simply based on residues modulo p of $d(u, v)$, we will unavoidably introduce many 4-cycles. For this reason, a resolution of this conjecture even in the special case where F is the 5-cycle would be very interesting, as such an argument would likely overcome many of the difficulties involved in proving the full conjecture.

Another area for further research is determining the optimal constant c_F in the statement of Theorem 1.8. The following question was suggested by Nešetřil [110].

Question 4.24. *For each graph F containing an edge, what is the smallest value of c_F for which Theorem 1.8 holds? Can we take c_F bounded by a function of $\chi(F)$?*

Clearly c_F must be at least $\chi(F) - 1$ for all F, and at least $\chi(F)$ if F is not vertex-critical. Nešetřil [110] has conjectured that the optimal c_F is at most $\chi(F)$ for all F.

For triangles, it follows from results in [35, 134] that the optimal c_{K_3} is either 3 or 4 (in particular, the bound $\chi(F) - 1$ for vertex-critical graphs is not always attained). For larger cliques, Briański, Davies, and Walczak [31] showed that the optimal c_{K_m} is $O(m^3)$; as noted at the start of Section 4.2, the optimal c_F is $O(|V(F)|^9)$ in general.

Chapter 5

Planar inducibility

In this chapter we consider the problem of maximising the number of subgraphs of an n-vertex planar graph that are isomorphic to some fixed H, for various small graphs H. We will focus predominantly on the version of the problem in which the subgraphs isomorphic to H are required to be induced, but will also present some results in the setting where any subgraph isomorphic to H counts. Recall from Section 1.4 that we denote by $f_I(n, H)$ the maximum number of induced copies of H that can be contained in an n-vertex planar graph, and by $f(n, H)$ the maximum number of not necessarily induced copies.

Most existing work on this problem relates to the cases where the graph H whose number of copies we seek to maximise is either a path or a cycle. Alon and Caro's results in [3] are notable exceptions: they determined $f(n, H)$ exactly for all n in the cases where H is a complete bipartite graph or a maximal planar graph with no separating triangle. In the general case Huynh, Joret, and Wood [85] showed that $f(n, H) = \Theta(n^{F(H)})$ for all H, where $F(H)$ denotes the so-called *flap-number* of H, and Liu [100, Corollary 6.1] showed that for all graphs H, $f_I(n, H) = \Theta(f(n, H))$.

In the case where H is a cycle, $f(n, C_3)$, $f(n, C_4)$ (both [70]), and $f(n, C_5)$ [68] are known exactly for all n, and $f_I(n, C_4)$ [132], $f_I(n, C_5)$ ([61] and [132] independently), and $f_I(n, C_6)$ [131] are known exactly for large n. For longer even cycles, building on work of Cox and Martin [38, 39], the asymptotics of $f(n, C_{2k})$ for all $k \geqslant 3$ were determined in [101], and the same asymptotics hold for $f_I(n, C_{2k})$. For odd cycles, in [80] the asymptotics of $f(n, C_7)$ and $f(n, C_9)$ were determined and an asymptotic upper bound within a factor of 3/2 of the best known construction was established for all longer odd cycles.

In the case where H is a path, progress has been restricted to the non-induced version of the problem. We write P_n for the path on n vertices. The value of $f(n, H)$ is known exactly for all n in the case where $H = P_2$ (Euler's formula), $H = P_3$ [3], and $H = P_4$ [64], and is known asymptotically in the case $H = P_5$ [60] and $H = P_7$ [38]. Moreover, for longer paths of even length an upper bound within an absolute constant (independent of the path length) of the best known lower bound construction was given in [8].

Our main results on $f(n, H)$ and $f_I(n, H)$ concern the cases where H is the paw graph (that is, the four-vertex graph consisting of a triangle with a pendant edge), K_4^- (that is, the graph obtained from the complete graph on four vertices by deleting one edge), or a star. Recall that for graphs G_1 and G_2, the *graph join* of G_1 and G_2, denoted $G_1 + G_2$, is the graph consisting of a copy of G_1 and a copy of G_2 on disjoint vertex sets V_1 and V_2 respectively, where every vertex in V_1 is adjacent to every vertex in V_2. For a graph G we denote the complement of G by $\overline{G}$.

We begin with our result for induced paw graphs, the proof of which is the most involved of those in this chapter. We determine $f_I(n, H)$ exactly for sufficiently large n, and identify the unique extremal graph containing this many induced paws.

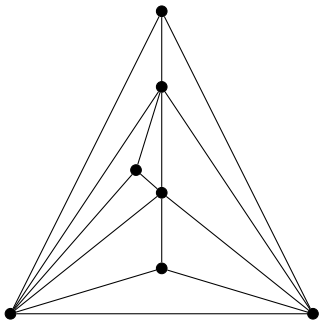
Theorem 5.1. *Let H denote the paw graph. For large n we have $f_I(n, H) = 2n^2 - 16n + 24$ and the unique n-vertex planar graph containing this many induced paws is $\overline{K_2} + C_{n-2}$.*

In the case of not necessarily induced paws, we are able to determine the extremal graphs for all n.

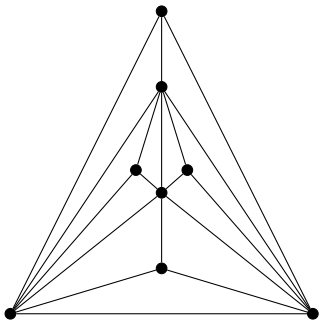
Theorem 5.2. *Let H denote the paw graph. For all $n \geqslant 4$ we have $f(n, H) = 4n^2 - 12n - 4$ and, unless $n \in \{7, 8\}$, the unique n-vertex planar graph containing this many paws is $K_2 + P_{n-2}$. When $n = 7$ or $n = 8$, there are exactly two extremal graphs: $K_2 + P_{n-2}$ and, respectively, the graphs A and B shown in Figure 5.1.*

Next we consider K_4^-. We begin with the induced case, where we identify the extremal graphs for all n.

Theorem 5.3. *For $n \geqslant 10$ we have $f_I(n, K_4^-) = \frac{1}{2}(n^2 - 3n - 4)$ and the unique n-vertex planar graph containing this many induced copies of K_4^- is $K_2 + P_{n-2}$. For*



(a) The exceptional 7-vertex graph, A



(b) The exceptional 8-vertex graph, B

Figure 5.1: The exceptional graphs for Theorems 5.2 and 5.4

other values of n we have

$$f_I(n, K_4^-) = \begin{cases} 1, & \text{if } n = 4 \\ 4, & \text{if } n = 5 \\ 12, & \text{if } n = 6 \\ 15, & \text{if } n = 7 \\ 19, & \text{if } n = 8 \\ 27, & \text{if } n = 9, \end{cases}$$

and the extremal graphs are as shown in Figure 5.2.

In the not necessarily induced case, we are again able to determine the extremal graphs for all n.

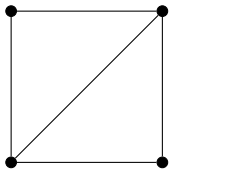
Theorem 5.4. *For all $n \geqslant 4$ we have $f(n, K_4^-) = \frac{1}{2}(n^2 + 9n - 40)$ and, unless $n \in \{7, 8\}$, the unique n-vertex planar graph containing this many copies of K_4^- is $K_2 + P_{n-2}$. When $n = 7$ or $n = 8$, there are exactly two extremal graphs: $K_2 + P_{n-2}$ and, respectively, the graphs A and B shown in Figure 5.1.*

Finally, we consider stars, denoting by S_k the star $K_{1,k}$ for $k \geqslant 2$. In the not necessarily induced case Alon and Caro [3] showed the following.

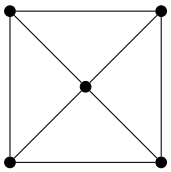
Theorem 5.5 ([3])**.** *For $k \geqslant 2$ and $n \geqslant 4$ we have*

$$f(n, S_k) = 2\binom{n-1}{k} + 2\binom{3}{k} + (n-4)\binom{4}{k}.$$

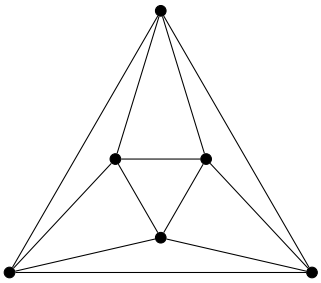
It follows from the proof given in [3] that for $n \geqslant k+1$, if $k \in \{2, 3, 4\}$ then the unique extremal graph is $K_2 + P_{n-2}$, and if $k \geqslant 5$ then the extremal graphs are exactly those of the form $K_2 + F$ where F is a spanning subgraph of P_{n-2}.



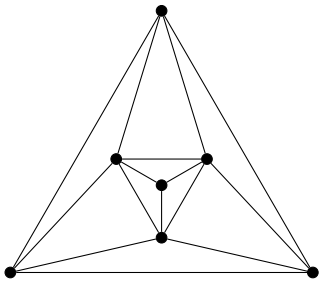
(a) The extremal 4-vertex graph



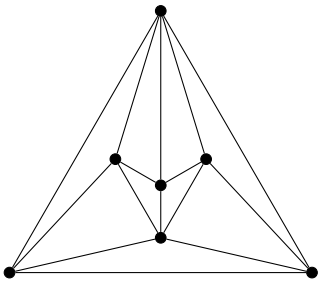
(b) The extremal 5-vertex graph



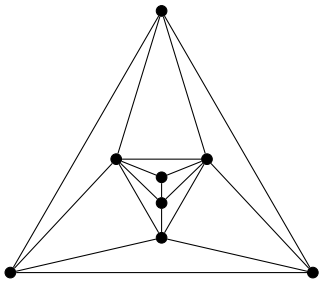
(c) The extremal 6-vertex graph



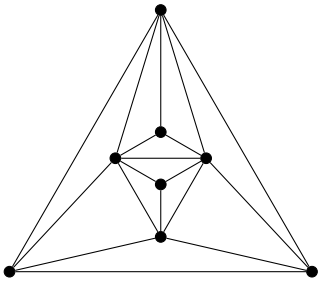
(d) The first extremal 7-vertex graph



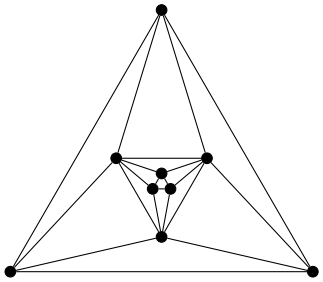
(e) The second extremal 7-vertex graph



(f) The first extremal 8-vertex graph



(g) The second extremal 8-vertex graph



(h) The extremal 9-vertex graph

Figure 5.2: The n-vertex planar graphs containing the most induced copies of K_4^- for small n

We investigate the induced case. For $k \in \{2,3,4\}$ we are able to determine the extremal graphs for all n, while for $k \geqslant 5$ our results only hold for large n. We begin with S_2, which is isomorphic to the path of length 2, P_3.

Theorem 5.6. *For each $n \geqslant 3$, we have $f_I(n, P_3) = (n-2)^2$ and the unique n-vertex planar graph achieving this maximum is $K_{2,n-2}$.*

For $k \geqslant 3$, there are, perhaps, three obvious candidates for an n-vertex planar graph containing the most induced copies of S_k: S_{n-1}, $K_{2,n-2}$, and $K_2 + \overline{K_{n-2}}$. The former contains $\binom{n-1}{k}$ induced copies of S_k, while the latter two each contain $2\binom{n-2}{k}$. Thus, we might expect that S_{n-1} is extremal for small n, with $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$ taking over as n increases. We show that this is the case when $k = 3$.

Theorem 5.7. *We have*

$$f_I(n, S_3) = \begin{cases} \binom{n-1}{3}, & \text{if } 4 \leqslant n \leqslant 7 \\ 2\binom{n-2}{3}, & \text{if } n \geqslant 7. \end{cases}$$

For $4 \leqslant n \leqslant 6$ the unique extremal graph is S_{n-1}, for $n \geqslant 8$ the extremal graphs are $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$, and for $n = 7$ all three of these graphs (and no others) are extremal.

Next, we show that the analogous result holds for S_4.

Theorem 5.8. *We have*

$$f_I(n, S_4) = \begin{cases} \binom{n-1}{4}, & \text{if } 5 \leqslant n \leqslant 9 \\ 2\binom{n-2}{4}, & \text{if } n \geqslant 9. \end{cases}$$

For $5 \leqslant n \leqslant 8$ the unique extremal graph is S_{n-1}, for $n \geqslant 10$ the extremal graphs are $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$, and for $n = 9$ all three of these graphs (and no others) are extremal.

We conjecture that this pattern extends to larger k but have been unable to prove it; our proofs of Theorems 5.7 and 5.8 rely on computer checks for small values of n, which become infeasible as k gets large. For each fixed k we can, however, prove that $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$ are the only extremal graphs when n is sufficiently large.

Theorem 5.9. *Fix $k \geqslant 3$. For sufficiently large n, we have $f_I(n, S_k) = 2\binom{n-2}{k}$ and the only n-vertex planar graphs containing this many induced copies of S_k are $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$.*

The proofs in this chapter contain many varied and ad hoc elements, but we now briefly outline the broad strokes of our arguments. We generally start by upper bounding the number of times that a vertex v of some degree d can appear in an (induced) copy of H in an n-vertex planar graph, or sometimes the number of times that v can appear as some *particular vertex* in such a copy of H. This can be combined with the lower bound on $f_I(n, H)$ or $f(n, H)$ provided by the graph(s) in the theorem statement to determine or bound the minimum degree of the extremal graphs. We then either induct on n (with base cases sometimes relying on a computer check), or, when this is not possible (for example, when our arguments in previous steps require n to be large), we apply a lemma stated and proved in Section 5.1.2 that translates the problem of determining $f_I(n, H)$ into the problem of characterising the planar graphs in which every 'small' set of vertices is in 'many' induced copies of H.

Chapter outline We begin in Section 5.1 by detailing some notation, definitions, and preliminary results. The proof of Theorem 5.1 on induced paws is tackled in Section 5.2. This is followed by the proof of Theorem 5.3 on induced copies of K_4^- in Section 5.3. Section 5.4 deals with the theorems concerning induced stars, with Theorems 5.6, 5.7, and 5.8 proved in Section 5.4.1 and Theorem 5.9 proved in Section 5.4.2. Our results on not necessarily induced subgraphs, Theorems 5.2 and 5.4, are proved in Section 5.5. In Section 5.6 we prove some corollaries of the main theorems, before closing in Section 5.7 with a discussion of $f_I(n, P_4)$.

5.1 Preliminaries

5.1.1 Notation and definitions

As above, for graphs G_1 and G_2, the *graph join* of G_1 and G_2, denoted $G_1 + G_2$, is the graph consisting of a copy of G_1 and a copy of G_2 on disjoint vertex sets V_1 and V_2 respectively, where every vertex in V_1 is adjacent to every vertex in V_2. We write K_n for the complete graph on n vertices, $K_{m,n}$ for the complete bipartite graph with parts of sizes m and n, C_n for the cycle of length n, P_n for the path on n vertices, and S_n for the star $K_{1,n}$. For a graph G, we write $\overline{G}$ for the *complement* of G, that is, the graph with the same vertex set as G in which vertices are neighbours exactly when they are not neighbours in G.

For a graph G and disjoint subsets $A, B \subseteq V(G)$, we will write $G[A, B]$ for the bipartite subgraph of G induced on parts A and B, that is, the graph with vertex set

$A \cup B$ and edge set $\big\{ab \in E(G) : a \in A, b \in B\big\}$. We write $\delta(G)$ for the minimum degree of a graph G. For $S \subseteq V(G)$ we write $\Gamma_G(S) \coloneqq \{u \in V(G) : \exists v \in S \text{ s.t. } uv \in E(G)\}$ for the neighbourhood of S in G, and when $S = \{v\}$ is a singleton we just write $\Gamma_G(v) = \Gamma_G(\{v\})$. In both cases we omit the subscript when the graph G is clear. For $S \subseteq V(G)$, $G - S$ is the subgraph of G induced on $V(G) \setminus S$, and again when $S = \{v\}$ is a singleton we simply write $G - v = G - \{v\}$.

5.1.2 Preliminary results

The first result of this subsection forms the cornerstone of our proofs of Theorems 5.1 and 5.9. It can be seen as reformulating the problem of determining $f_I(n, H)$ and the graphs attaining that maximum into the problem of characterising the planar graphs in which every 'small' set of vertices is in 'many' induced copies of H. This approach was used on an ad hoc basis in [61, 131, 132]. Since we will want to use it in more than one setting, we will formalise it in a lemma, the proof of which is a generalisation of the proof of [61, Theorem 3].

Before stating the result we give the following definition. Let H be a planar graph, and for each n let $\mathcal{G}_n$ be a non-empty family of n-vertex planar graphs that all contain the same number of induced copies of H. Write $h(n)$ for the number of such copies in each graph in $\mathcal{G}_n$.

Definition 5.10. For each $s \geqslant 1$, we say that H and $(\mathcal{G}_n)$ satisfy *condition* $(\star)_s$ if for sufficiently large n, every n-vertex planar graph that

(a) contains at least $h(n)$ induced copies of H, and

(b) has the property that for all non-empty $S \subseteq V(G)$ with $|S| \leqslant s$ there are at least $h(n) - h(n - |S|)$ induced copies of H intersecting S

belongs to $\mathcal{G}_n$.

We now state the key lemma.

Lemma 5.11. *Let H, $(\mathcal{G}_n)$, and $h(n)$ be as above and suppose there exists some $s \geqslant 1$ such that they satisfy condition $(\star)_s$. Then for large n, we have $f_I(n, H) = h(n)$ and the n-vertex planar graphs achieving this maximum are exactly the graphs in $\mathcal{G}_n$.*

Proof. First note that since each $\mathcal{G}_n$ is non-empty we have $f_I(n, H) \geqslant h(n)$ for all n. Let n_0 be sufficiently large in the sense of condition $(\star)_s$, then let $n_1 = n_0 + sn_0^{|V(H)|}$. Suppose for a contradiction that there exists $n \geqslant n_1$ such that $f_I(n, H) > h(n)$, and

for such an n let G be an n-vertex planar graph containing at least $h(n)+1$ induced copies of H. Then $G \notin \mathcal{G}_n$, so by condition $(\star)_s$ there is some non-empty $S \subseteq V(G)$ with $|S| \leqslant s$ such that there are fewer than $h(n) - h(n-|S|)$ induced copies of H in G containing a vertex in S. Deleting the vertices in S from G yields a graph on $n-|S|$ vertices containing at least $h(n-|S|)+2$ induced copies of H. Repeat this process until we first obtain a graph on fewer than n_0 vertices, say on n_2 vertices. The n_2-vertex planar graph we obtain contains at least $n_0^{|V(H)|}$ induced copies of H, which is impossible. This gives the desired contradiction, and we deduce that for $n \geqslant n_1$ we have $f_I(n,H) = h(n)$.

It remains to show that for large enough n, every n-vertex planar graph containing $h(n)$ induced copies of H is in $\mathcal{G}_n$. For $n \geqslant n_1 + s$, every n-vertex planar graph containing $f_I(n,H)$ induced copies of H has the property that for all non-empty $S \subseteq V(G)$ with $|S| \leqslant s$ there are at least $h(n) - h(n-|S|)$ induced copies of H containing a vertex in S (else we could delete the vertices in S to obtain a graph contradicting the conclusion of the previous paragraph), so every such graph is in $\mathcal{G}_n$ by condition $(\star)_s$, as required. □

Our second preliminary result is a weakened version of a lemma used by Alon and Caro in [3]. We include a short proof from that paper.

Lemma 5.12 ([3])**.** *Let G be a plane graph. Suppose that $x, y, z \in V(G)$ are the vertices of a triangle in G, and let S be the set of vertices in the interior of the bounded region formed by the triangle xyz. Suppose that S is non-empty and that every vertex in S is adjacent to x. Then there exists a vertex in S which has degree at most 3 in G.*

Proof. If $|S| = 1$, then the result is clear. Otherwise, note that $G' = G[S \cup \{y,z\}]$ is an outerplanar graph on at least four vertices. It is well known and easy to see that such graphs contain two non-adjacent vertices each of degree at most 2. In particular some vertex in S has degree at most 2 in G', and hence degree at most 3 in G. □

We will also make frequent use of various basic facts about planar graphs, such as that every maximal n-vertex planar graph has exactly $3n-6$ edges for $n \geqslant 3$, from which it follows that planar graphs have minimum degree at most 5. We will also regularly use the easy direction of Kuratowski's theorem [98], and our appeals to this result will usually be "by planarity".

Theorem 5.13 (Kuratowski's theorem, [98])**.** *A graph is planar if and only if it does not contain a subgraph isomorphic to a subdivision of either K_5 or $K_{3,3}$.*

5.1.3 A note on computer checks

We use computer checks in the proofs of Theorems 5.3, 5.7, and 5.8 to verify small cases. For this, we used lists of connected planar graphs sourced from The House of Graphs [37]. Note that since every graph H whose number of copies we seek to maximise in these theorems is connected, every planar n-vertex graph containing $f(n, H)$ copies of H or $f_I(n, H)$ induced copies of H is a subgraph of a *connected* n-vertex planar graph with the same property. Hence, to find all n-vertex extremal graphs it is sufficient to find the extremal *connected* graphs, and then check whether any disconnected subgraphs of these are also extremal. In our applications, it will be clear that no such disconnected subgraphs exist.

5.2 Induced paws

This section is devoted to proving Theorem 5.1 via an application of Lemma 5.11 with $s = 2$. In the notation of that lemma, we take $\mathcal{G}_n = \{\overline{K_2} + C_{n-2}\}$ which we will show gives $h(n) = 2n^2 - 16n + 24$, and hence $h(n) - h(n-1) = 4n - 18$ and $h(n) - h(n-2) = 8n - 40$.

We start with a lemma which gives an upper bound on the number of times that a vertex of degree 3 can appear as the pendant vertex in an induced paw in a planar graph, and which gives a necessary condition for that upper bound to be attained. This serves as the base case for the more general Lemma 5.15, as well as helping us to count induced paws in Lemma 5.18.

Lemma 5.14. *Let $n \geqslant 5$ and let G be an n-vertex planar graph. If $v \in V(G)$ has degree 3, then G contains at most $2n - 10$ induced paws in which v is the pendant vertex. Moreover, if G contains $2n - 10$ such paws, then either*

(a) adding any edges between vertices of $\Gamma(v)$ which are not already present in G yields a maximal planar graph in which some vertex other than v is adjacent to every neighbour of v, or

(b) there exists $u \in \Gamma(v)$ with $d(u) \leqslant 3$.

Proof. Suppose that $v \in V(G)$ has degree 3, and let u_1, u_2, and u_3 be the neighbours of v. We wish to determine the number of induced paws in G in which v is the pendant vertex. The addition of any of the edges u_1u_2, u_2u_3, or u_1u_3 to G (if they are not already present) does not affect the number of such paws, nor does it affect the planarity of the graph. We may therefore assume that all three of these edges are present. To prove the lemma, it is sufficient to show that there are at most $2n - 10$ induced paws in G in which v is the pendant vertex, with equality only if v has a neighbour of degree 3 or G is a maximal planar graph in which some vertex other than v is adjacent to every neighbour of v.

Let A and B be the sets of vertices in $V(G) \setminus (\{v\} \cup \Gamma(v))$ which are adjacent to at least two neighbours of v and exactly one neighbour of v respectively. Let $|A| = a$ and $|B| = b$, and let $e(A)$, $e(B)$, and $e(A, B)$ denote the number of edges in $G[A]$, $G[B]$, and $G[A, B]$ respectively. Each edge in $G[A]$ is in at most two induced paws in which v is the pendant vertex (since no two vertices in A are neighbours of every vertex in $\Gamma(v)$ by planarity), and each edge in $G[B]$ and $G[A, B]$ is in at most one. Thus, the number of induced paws in G in which v is the pendant vertex is at most $2e(A) + e(B) + e(A, B)$. The remainder of the proof is split into the following cases.

Case 1: Suppose that G contains a vertex other than v which is adjacent to all of u_1, u_2, and u_3. Note that such a vertex is unique by the planarity of G. Consider the graph $G[\{v\} \cup \Gamma(v) \cup A \cup B]$, which contains $a + b + 4$ vertices and thus at most $3(a + b + 4) - 6 = 3a + 3b + 6$ edges. There are six edges in $G[\{v\} \cup \Gamma(v)]$, $2a + 1$ edges in $G[A, \Gamma(v)]$, and b edges in $G[B, \Gamma(v)]$, so we have

$$3a + 3b + 6 \geqslant 6 + (2a + 1) + b + e(A) + e(B) + e(A, B).$$

Rearranging, we obtain

$$e(A) + e(B) + e(A, B) \leqslant a + 2b - 1. \tag{5.1}$$

Now consider $G[\{v\} \cup \Gamma(v) \cup A]$. This graph contains $a + 4$ vertices and hence at most $3a + 6$ edges. Counting edges as above gives $3a + 6 \geqslant 6 + (2a + 1) + e(A)$, which rearranges to $e(A) \leqslant a - 1$. Adding this to inequality (5.1) gives $2e(A) + e(B) + e(A, B) \leqslant 2a + 2b - 2$. Since $a + b \leqslant n - 4$, this implies that $2e(A) + e(B) + e(A, B) \leqslant 2n - 10$, from which we deduce that there are at most $2n - 10$ induced paws in which v is the pendant vertex. If equality holds here, then equality holds throughout so in particular $a + b = n - 4$. In this case $G[\{v\} \cup \Gamma(v) \cup A \cup B] = G$ and this graph contains $3n - 6$ edges, so G is a maximal planar graph.

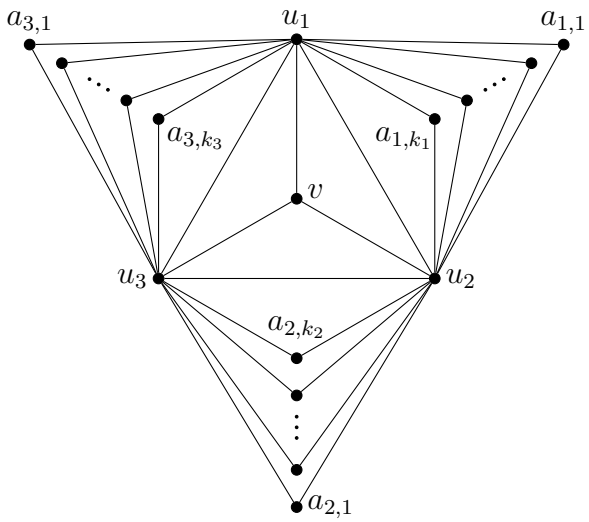


Figure 5.3: An illustration of part of the graph G' in Case 2(i)

Case 2(i): Suppose that v is the only vertex adjacent to all of u_1, u_2, and u_3, but that each pair from u_1, u_2, and u_3 have a common neighbour in A. We will show that in this case G contains at most $2n-11$ induced paws in which v is the pendant vertex. Let $G' = G[\{v\} \cup \Gamma(v) \cup A]$. Draw this graph and consider the induced drawing of the triangle $u_1u_2u_3$. We may assume that v is the only vertex in the interior of the bounded region formed by this triangle.

For each $i \in \{1,2,3\}$, let $a_{i,1}, a_{i,2}, \ldots, a_{i,k_i}$ be the elements of $(\Gamma_{G'}(u_i) \cap \Gamma_{G'}(u_{i+1})) \setminus \{v, u_{i+2}\}$ labelled in natural order according to the drawing of G' (see Figure 5.3), where addition in the subscripts of the u_i is taken modulo 3. Since no vertex in A is in the common neighbourhood of all of u_1, u_2, and u_3, the sets $\{a_{i,1}, a_{i,2}, \ldots, a_{i,k_i}\}$ for $i \in \{1,2,3\}$ are disjoint, and since each pair from u_1, u_2, and u_3 have a common neighbour in A, we have $k_i \geqslant 1$ for all i.

Let $\ell \in \{0,1,2,3\}$ be the number of edges present amongst $a_{1,1}$, $a_{2,1}$, and $a_{3,1}$ in G (and hence also in G'). Observe that any edges not present amongst these vertices can be added to G' while maintaining planarity. It follows that $e(G') \leqslant 3(a+4)-6-(3-\ell) = 3a+3+\ell$. We also have $e(G') = 6+2a+e(A)$, so we deduce that

$$e(A) \leqslant a - 3 + \ell. \tag{5.2}$$

Similarly to in Case 1, $G[\{v\} \cup \Gamma(v) \cup A \cup B]$ has at most $3a+3b+6$ edges, and this must be at least $6+2a+b+e(A)+e(B)+e(A,B)$. Rearranging yields $e(A)+e(B)+e(A,B) \leqslant a+2b$. Adding this to inequality (5.2), we obtain $2e(A)+e(B)+e(A,B) \leqslant 2a+2b-3+\ell$. Since $a+b \leqslant n-4$, this is at most $2n-11+\ell$.

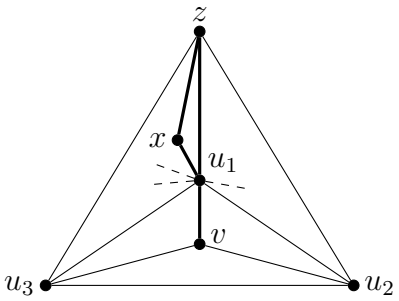


Figure 5.4: The extra paw induced on $\{v, u_1, x, z\}$ that gives a contradiction in Case 2(ii). The dashed lines represent other edges incident to u_1 that may or may not be present

Each edge amongst $a_{1,1}$, $a_{2,1}$, and $a_{3,1}$ in G is an edge of $G[A]$ that is in only one induced paw in which v is the pendant vertex. It follows that the total number of induced paws in G in which v is the pendant vertex is at most $2e(A)+e(B)+e(A,B)-\ell$. We have shown that this is at most $2n-11$, as claimed.

Case 2(ii): Now suppose simply that no vertex other than v is adjacent to all of u_1, u_2, and u_3. Fix a drawing of G and note that we may assume that v is the only vertex in the interior of the bounded region formed by the triangle $u_1u_2u_3$ in the drawing. Construct a plane triangulation F from G by adding arbitrary edges one-by-one to the unbounded region formed by the triangle $u_1u_2u_3$ in the drawing of G. Every induced paw in G in which v is the pendant vertex also appears as an induced subgraph of F, so in particular there are at least as many such paws in F as there are in G. Since $n \geqslant 5$, there is at least one vertex in the unbounded region formed by the triangle $u_1u_2u_3$, so each pair from u_1, u_2, and u_3 have a common neighbour in this region in F. The triangulation therefore satisfies the conditions of either Case 1 or Case 2(i) and hence contains at most $2n-10$ induced paws in which v is the pendant vertex. It follows that the same is true of G.

If we have equality in G, then any graph F constructed from G as above is of the form considered in Case 1, so there is a step in the construction process at which, without loss of generality, we add an edge from u_1 to a vertex $z \in V(G) \setminus (\{v\} \cup \Gamma(v))$ which is adjacent to u_2 and u_3 in the current graph. If $d_G(u_1) \geqslant 4$, then there exists $x \in \Gamma_G(u_1) \setminus \{u_2, u_3, v\}$ such that just before we add the edge u_1z, some face of the graph is incident to u_1, z, and x. If x and z are not neighbours in this graph, then after we have added the edge u_1z we can add the edge xz, and then continue with an arbitrary construction of a graph F. If x and z are already neighbours, then we just add the edge u_1z and continue an arbitrary construction of F from there. In

either case, the resulting graph F contains exactly $2n-10$ induced paws in which v is the pendant vertex, one of which is induced on $\{v, u_1, x, z\}$ (see Figure 5.4). This copy is not present in G, so G contains fewer than $2n-10$ induced paws in which v is the pendant vertex. This is a contradiction, so if G contains $2n-10$ induced paws in which v is the pendant vertex, then some neighbour of v has degree 3 in G, as required. □

We now generalise Lemma 5.14 as follows.

Lemma 5.15. *Let $k \geqslant 3$ and $n \geqslant k+2$. Let G be an n-vertex planar graph and suppose that $v \in V(G)$ has degree k. Then the number of induced paws in G in which v is the pendant vertex is at most $2n-2k-4$. Moreover, if equality holds, then either*

(a) given any drawing of G the addition of any edges between consecutive neighbours of v which are not already present yields a maximal planar graph in which $\Gamma(v)$ induces a cycle, or

(b) there exists $u \in \Gamma(v)$ with $d(u) \leqslant 3$.

Proof. We will prove the lemma by induction on k. The case $k = 3$ follows from Lemma 5.14, so let $k \geqslant 4$ and assume that the result holds for all smaller values of k. Let n, G, and v be as in the statement of the lemma. Observe first of all that the addition to G of any edges which are not incident to v cannot destroy any induced paws in which v is the pendant vertex. Fix a drawing of G and label the neighbours of v as $u_1, \ldots, u_k$ in the order they appear around v. By the above observation, we may assume that the cycle $u_1u_2 \ldots u_k$ is present in G.

Construct a plane graph G_0 from G by arbitrarily adding edges between vertices of $\Gamma_G(v)$ to the drawing one-by-one while preserving planarity, until no further such edges can be added. Note that G_0 contains at least as many induced paws in which v is the pendant vertex as G does (in fact, the copies are the same in each). If there exists $i \in [k]$ such that $d_{G_0}(u_i) \leqslant 3$, then there are no induced paws in G_0 which contain u_i and in which v is the pendant vertex. By the induction hypothesis, $G_0 - u_i$ contains at most $2(n-1)-2(k-1)-4 = 2n-2k-4$ induced paws in which v is the pendant vertex, so the same is true of G_0, and hence of G. We may therefore assume that $d_{G_0}(u_i) \geqslant 4$ for all $i \in [k]$. The following claim shows that we may further assume that $G_0[\Gamma(v)]$ is exactly the k-cycle $u_1 \ldots u_k$ (or, equivalently, that $G = G_0$).

Claim 5.16. *If $G_0[\Gamma(v)]$ is not a k-cycle, then G_0 contains at most $2n - 2k - 6$ induced paws in which v is the pendant vertex.*

Proof. Without loss of generality, assume that G_0 contains the edge u_1u_i for some $i \in \{3, \ldots, k-1\}$. In the drawing of G_0, the triangle vu_1u_i splits the plane into two regions, both of whose interiors contain neighbours of v. Let A denote the set of vertices in triangle vu_1u_i and its interior and B denote the set of vertices in triangle vu_1u_i and its exterior. The number of induced paws in G_0 in which v is the pendant vertex is the number of such paws in $G_0[A]$ plus the number of such paws in $G_0[B]$. By Lemma 5.12 and the fact that $d_{G_0}(u_i) \geqslant 4$ for all $i \in [k]$, each of A and B contain a vertex not in $\{v\} \cup \Gamma(v)$, so we can apply the induction hypothesis to $G_0[A]$ and $G_0[B]$. It follows that the number of induced paws in G_0 in which v is the pendant vertex is at most

$$\begin{aligned} 2(|A| + |B|) - 2(d_{G_0[A]}(v) + d_{G_0[A]}(v)) - 8 &= 2(n+3) - 2(k+2) - 8 \\ &= 2n - 2k - 6, \end{aligned}$$

as required. □

It follows from this claim that we may assume that $G = G_0$. In the drawing of G, the cycle $u_1u_2 \ldots u_k$ splits the plane into two regions, one of which, say the bounded one, we may assume contains only v among the vertices of G. Construct a plane graph G_1 from G by repeatedly arbitrarily adding sets of edges S to the unbounded region formed by the cycle $u_1u_2 \ldots u_k$ while preserving planarity, with the condition that after the addition of each set S the number of induced paws in the graph in which v is the pendant vertex must have increased. Continue until no further sets S can be added. To prove the lemma it is sufficient to show that G_1 contains at most $2n - 2k - 4$ such paws, with equality only if G_1 is a plane triangulation.

In G_1, u_2 is adjacent to exactly u_1 and u_3 among the neighbours of v. Let $d_{G_1}(u_2) = d + 3 \geqslant 4$, and let $x_1, \ldots, x_d$ be the neighbours of u_2 other than u_1, u_3, and v, where we label these neighbours in the order they appear around u_2 in the drawing of G_1, with x_1 next to u_1 and x_d next to u_3 (see Figure 5.5). Let $X = \{x_1, \ldots, x_d\}$. By the construction of G_1 we may assume that x_i and x_{i+1} are neighbours for all $i \in [d-1]$.

We now show that u_1 and x_1 are neighbours. Suppose for a contradiction that they are not and consider the face of G_1 that is incident to u_1, u_2, and x_1. This face has at least four vertices on its boundary: let z be a vertex on the boundary other than u_1, u_2, and x_1. Clearly $z \neq v$. The edge u_1z is either already present in G_1 or can

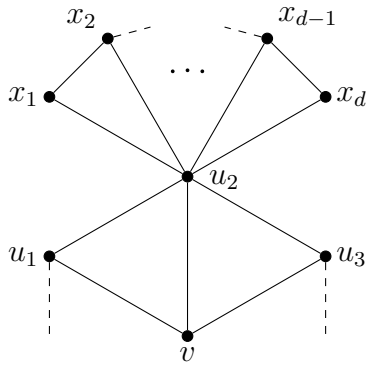


Figure 5.5: An illustration of part of the drawing of G_1

be added while preserving planarity. If $z \in \{u_3, \ldots, u_{k-1}\}$, then this contradicts the construction of G_0. Similarly, the edge u_2z is either already present in G_1 or can be added while preserving planarity, and if $z = u_k$ then this contradicts the construction of G_0. It follows that $z \in V(G_1) \setminus (\{v, x_1\} \cup \Gamma(v))$. Any of the edges among u_1, x_1, and z that are not already present in G_1 can be added simultaneously while preserving planarity, but this introduces a new induced paw on vertex set $\{v, u_1, z, x_1\}$ in which v is the pendant vertex (and does not destroy any such paws). This contradicts the construction of G_1. Hence, u_1 and x_1 are neighbours in G_1, and similarly for u_3 and x_d.

Let N be the number of induced paws in G_1 in which v is the pendant vertex. We may assume that in the drawing of G_1 the cycle $x_1 \ldots x_d u_3 v u_1$ splits the plane into two regions, one of which contains only u_2 among the vertices of G_1. The number of induced paws in G_1 which contain u_2 and in which v is the pendant vertex is $e(G_1[X])$. Let F_0 be the plane graph $G_1 - u_2$. Then F_0 has $n-1$ vertices and contains exactly $N - e(G_1[X])$ induced paws in which v is the pendant vertex. Let P be the face of F_0 which contains u_2 in the drawing of G_1, so that P is bounded by the cycle $x_1 \ldots x_d u_3 v u_1$. Define $S_1 = \{x \in X : x \in \Gamma_{F_0}(u_1)\} = \{x \in X : x \in \Gamma_{G_1}(u_1)\}$ and $S_2 = X \setminus S_1$, and let $\sigma = |S_2|$, so that $|S_1| = d - \sigma$.

Construct a plane graph F_1 from F_0 by adding the edges $u_1 s$ for $s \in S_2$ to face P. Then F_1 contains exactly $N - e(G_1[X]) + e(G_1[S_2]) + e(G_1[S_1, S_2]) = N - e(G_1[S_1])$ induced paws in which v is the pendant vertex. Let $P_0, \ldots, P_\sigma$ be the faces of F_1 into which face P of F_0 is split. Let $1 < i_1 < i_2 < \cdots < i_\sigma \leqslant d$ be such that $S_2 = \{x_{i_j} : 1 \leqslant j \leqslant \sigma\}$, and let $i_0 = 1$. Then for $0 \leqslant j \leqslant \sigma - 1$ face P_j has boundary given by the cycle $u_1 x_{i_j} x_{i_j+1} \ldots x_{i_{j+1}}$ and face P_σ has boundary given by the cycle $u_1 x_{i_\sigma} x_{i_\sigma+1} \ldots x_d u_3 v$.

Claim 5.17. *If $i_\sigma \leqslant d-1$, then for each $i_\sigma \leqslant t \leqslant d-1$, u_3 and x_t are not neighbours in F_1.*

Proof. Suppose for a contradiction that $x_t \in \Gamma_{F_1}(u_3)$ for some $i_\sigma \leqslant t \leqslant d-1$. Add the edges vx_t and vx_d to face P_σ of F_1 and let the resulting plane graph be called F'. We will show that u_1, u_3, v, x_t, and x_d are the vertices of degree 4 of a subdivision of K_5 in F', which would contradict the planarity of F'. First note that v is adjacent in F' to the four other vertices, u_3 is adjacent to x_d and to x_t by assumption, and u_1 is adjacent to x_t and x_d by the construction of F_1. The paths $x_t x_{t+1} \dots x_d$ and $u_3 u_4 \dots u_k u_1$ are vertex-disjoint paths in F' from x_t to x_d and from u_3 to u_1 respectively. This completes the proof of the claim. □

Construct a plane graph F_2 from F_1 as follows. First, for each $0 \leqslant j \leqslant \sigma-1$ triangulate the face P_j by adding the $i_{j+1} - i_j - 1$ edges $x_{i_j}x_{i_j+2}, x_{i_j}x_{i_j+3}, \dots, x_{i_j}x_{i_{j+1}}$. It is straightforward to see by planarity that none of these edges are present in F_1. Indeed, the edges $u_1x_{i_j+1}, u_1x_{i_j+2}, \dots, u_1x_{i_{j+1}-1}$ are all present in F_1 and are drawn outside the face P_j. Next, if $i_\sigma \leqslant d-1$, then add to P_σ the $d-i_\sigma$ edges $u_3x_{i_\sigma}, u_3x_{i_\sigma+1}, \dots, u_3x_{d-1}$, none of which are present in F_1 by Claim 5.17. Observe that this is not a triangulation of P_σ, but that if we were to also add the edge u_1u_3 (as we shall do in a moment), then it would be. This completes the construction of F_2.

In total, we add exactly $d-i_0-\sigma = d-1-\sigma$ edges in this construction process. Each of these additional edges introduces to the graph a new induced paw in which v is the pendant vertex, and no induced paws in which v is the pendant vertex are destroyed in the construction process. It follows that F_2 contains at least $N-e(G_1[S_1])+d-1-\sigma$ induced paws in which v is the pendant vertex. In G_1, S_1 is contained in the common neighbourhood of u_1 and u_2, and u_1 and u_2 are neighbours in this graph, so by the planarity of G_1, $G_1[S_1]$ is acyclic. We have $|S_1| \geqslant 1$ since $x_1 \in S_1$, so $e(G_1[S_1]) \leqslant |S_1| - 1 = d - \sigma - 1$. Therefore, F_2 contains at least as many induced paws in which v is the pendant vertex as G_1 does.

The graph F_2 is an $(n-1)$-vertex plane graph in which v has degree $k-1$, so by the induction hypothesis it contains at most $2(n-1) - 2(k-1) - 4 = 2n - 2k - 4$ induced paws in which v is the pendant vertex. Hence, the same is true of G_1. Since u_1 and u_3 are not neighbours in F_2, we can define a plane graph F_3 by adding the edge u_1u_3 to the face of F_2 incident to u_1, u_3, and v. If G_1 contains $2n-2k-4$ induced paws in which v is the pendant vertex, then so does F_3. In this case, by the induction hypothesis, F_3 is a plane triangulation. It is straightforward to see, using the

observations made during the construction of F_2, that if F_3 is a plane triangulation, then the same is true of G_1. This completes the proof of Lemma 5.15. □

Next, we show that if an n-vertex planar graph G satisfies part (b) of condition $(\star)_2$ with $h(n) = 2n^2 - 16n + 24$, then G has minimum degree at least 4.

Lemma 5.18. *Let $n \geqslant 13$ and let G be an n-vertex planar graph in which*

(i) every vertex of G is in at least $4n - 18$ induced paws, and

(ii) for every pair of vertices in G there are at least $8n - 40$ induced paws in G containing at least one vertex from the pair.

Then $\delta(G) \geqslant 4$.

Proof. First, since $4n - 18 > 0$, the fact that no vertex has degree 0 is trivial. If there exists $v \in V(G)$ with $d(v) = 1$, then any induced paw containing v is determined by the edge of the copy which does not contain v, so very crudely there are at most $3n - 6 < 4n - 18$ induced paws containing v.

Suppose $v \in V(G)$ has $d(v) = 2$. Let u_1 and u_2 be the neighbours of v, let $A = (\Gamma(u_1) \cap \Gamma(u_2)) \setminus \{v\}$, and let $B = (\Gamma(u_1) \triangle \Gamma(u_2)) \setminus \{u_1, u_2\}$. Write $a = |A|$ and $b = |B|$. The number of induced paws in which v is the pendant vertex is at most $2e(G[A]) + e(G[B]) + e(G[A, B])$. Proceeding as in the proof of Lemma 5.14, considering the edges of $G[\{v\} \cup \Gamma(v) \cup A \cup B]$ gives the inequality $3(a + b + 3) - 6 \geqslant 2 + 2a + b + e(G[A]) + e(G[B]) + e(G[A, B])$, so $e(G[A]) + e(G[B]) + e(G[A, B]) \leqslant a + 2b + 1$. Considering the edges of $G[\{v\} \cup \Gamma(v) \cup A]$ gives $3(a + 3) - 6 \geqslant 2 + 2a + e(G[A])$, which rearranges to $e(G[A]) \leqslant a + 1$. Adding the two inequalities shows that the number of induced paws in which v is the pendant vertex is at most $2(a + b) + 2$. We have $a + b \leqslant n - 3$, so this is at most $2n - 4$. There are at most $n - 3$ induced paws in which v is one of the vertices of degree 2, and none in which it is the vertex of degree 3. Hence the total number of induced paws containing v is at most $3n - 7 < 4n - 18$ so we have the desired contradiction.

Finally suppose that $v \in V(G)$ has $d(v) = 3$. If there is an induced paw in which v is the vertex of degree 3, then there are at most $n - 4$ induced paws in which v is one of the vertices of degree 2. By Lemma 5.14, there are at most $2n - 10$ induced paws in which v is the pendant vertex, so in this case there are at most $1 + (n - 4) + (2n - 10) = 3n - 13 < 4n - 18$ induced paws containing v. Hence there is no induced paw in which v is the vertex of degree 3.

Let $w \in V(G) \setminus (\{v\} \cup \Gamma(v))$. If w has exactly one or two common neighbours with v, then there are at most two induced paws in G which contain w and in which v is one of the vertices of degree 2. If w has exactly zero or three common neighbours with v, then there are no such paws in G. It follows that there are at most $2(n-4) = 2n-8$ induced paws in which v is one of the vertices of degree 2. Combining this with Lemma 5.14 gives an upper bound of $4n-18$ on the number of induced paws which contain v. Equality holds here by assumption, so equality holds in Lemma 5.14 and every vertex in $V(G) \setminus (\{v\} \cup \Gamma(v))$ has exactly one or two common neighbours with v. Consequently, by Lemma 5.14 there exists $u \in \Gamma(v)$ with $d(u) \leqslant 3$. If we delete v from G, then u has degree at most 2 in the $(n-1)$-vertex graph produced, so by the above it is contained in at most $3(n-1) - 7 = 3n - 10$ induced paws in that graph. It follows that the number of induced paws containing u or v is at most $4n - 18 + 3n - 10 = 7n - 28 < 8n - 40$, which gives the required contradiction. □

In our final lemma, we show that if an n-vertex planar graph G has minimum degree at least 4 and contains at least $h(n) = 2n^2 - 16n + 24$ induced paws, then G is the extremal graph from Theorem 5.1, $\overline{K_2} + C_{n-2}$.

Lemma 5.19. *Let n be large and let G be an n-vertex planar graph with $\delta(G) \geqslant 4$. If G contains at least $2n^2 - 16n + 24$ induced paws, then $G = \overline{K_2} + C_{n-2}$.*

Proof. Let n and G be as in the statement of the lemma and suppose that G contains at least $2n^2 - 16n + 24$ induced paws. We begin with the following claim.

Claim 5.20. $e(G) \geqslant 3n - 31$

Proof. There exists $v \in V(G)$ such that G contains at least $\lceil (2n^2 - 16n + 24)/n \rceil = 2n - 15$ induced paws in which v is the pendant vertex. By Lemma 5.15, we have $d(v) = d \in \{4, 5\}$. We now apply the ideas used in the proof of Lemma 5.14 to count the number of induced paws in which v is the pendant vertex. Let A and B be the sets of vertices in $V(G) \setminus (\{v\} \cup \Gamma(v))$ which are adjacent to at least two neighbours of v and exactly one neighbour of v respectively, and write $a = |A|$ and $b = |B|$. As in the proof of Lemma 5.14, the number of times that v appears as the pendant vertex in an induced paw is at most $2e(A) + e(B) + e(A, B)$.

Counting the edges in $G[\{v\} \cup \Gamma(v) \cup A \cup B]$ yields

$$e(\{v\} \cup \Gamma(v) \cup A \cup B) \geqslant d + 2a + b + e(A) + e(B) + e(A, B),$$

so

$$3(a+b+d+1)-6 \geqslant d+2a+b+e(A)+e(B)+e(A,B), \tag{5.3}$$

which rearranges to

$$e(A)+e(B)+e(A,B) \leqslant a+2b+2d-3. \tag{5.4}$$

Counting the edges in $G[\{v\} \cup \Gamma(v) \cup A]$ we obtain $3(a+d+1)-6 \geqslant d+2a+e(A)$, which rearranges to $e(A) \leqslant a+2d-3$. Adding this to inequality (5.4) yields $2e(A)+e(B)+e(A,B) \leqslant 2a+2b+4d-6$. Thus we have $2n-15 \leqslant 2(a+b)+4d-6$, which rearranges to $a+b \geqslant n-9/2-2d$, so in fact we have $a+b \geqslant n-2d-4$.

Suppose for a contradiction that $e(G) \leqslant 3n-32$. Then we may replace the left-hand side of inequality (5.3) with $3n-32$. Following the argument through in the same way as before yields $2n-15 \leqslant 2e(A)+e(B)+e(A,B) \leqslant 3n-a-b+d-35$. Now substituting in $a+b \geqslant n-2d-4$ gives $2n-15 \leqslant 2n+3d-31$. Since $d \leqslant 5$, this is a contradiction. □

By Lemma 5.15, the total number of induced paws in G is at most

$$\sum_{v \in V(G)} (2n-2d(v)-4) = 2n^2-4e(G)-4n \leqslant 2n^2-4(3n-31)-4n = 2n^2-16n+124.$$

There are at least $2n^2-16n+24$ induced paws in G by assumption, so at most 100 vertices $v \in V(G)$ do not appear as the pendant vertex in the maximum $2n-2d(v)-4$ induced paws. Since n is large and G is planar, there therefore exists $v \in V(G)$ with $d(v) \leqslant 6$ such that there are $2n-2d(v)-4$ induced paws in G in which v is the pendant vertex. Fix a drawing of G, then by Lemma 5.15 any faces of the drawing that are not triangles are incident to v, and G is at most six edges short of being a plane triangulation. There are at most six faces of the drawing which are incident to v, so there are only a bounded number of vertices of G incident to a non-triangular face of the drawing. It follows that there exists $u \in V(G)$ such that G contains $2n-2d(u)-4$ induced paws in which u is the pendant vertex, and such that every face incident to u in the drawing of G is a triangle. Lemma 5.15 now implies that G is a plane triangulation.

Applying Lemma 5.15 again we find that the total number of induced paws in G is at most

$$\sum_{v \in V(G)} 2n-2d(v)-4 = 2n^2-4e(G)-4n = 2n^2-4(3n-6)-4n = 2n^2-16n+24,$$

so in fact G contains exactly this many induced paws, and for every $v \in V(G)$ there are exactly $2n - 2d(v) - 4$ induced paws in which v the pendant vertex. Hence, by Lemma 5.15, the neighbourhood of every vertex of G induces exactly a cycle.

Now let $v \in V(G)$ and consider the induced paws in which v is the vertex of degree 3. The number of such copies is the number of pairs consisting of an edge of $G[\Gamma(v)]$ and a vertex in $\Gamma(v)$ which is not adjacent to either endpoint of the edge. Since $G[\Gamma(v)]$ is a cycle, there are exactly $d(v)(d(v) - 4)$ such induced paws in G. Hence

$$2n^2 - 16n + 24 = \sum_{v \in V(G)} (d(v)^2) - 8e(G) = \sum_{v \in V(G)} (d(v)^2) - 24n + 48.$$

This rearranges to $\sum_{v \in V(G)} d(v)^2 = 2n^2 + 8n - 24$. Every vertex in G has degree between 4 and $n - 2$ inclusive (if $v \in V(G)$ has degree $n - 1$ then it is the pendant vertex in exactly $2n - 2(n - 1) - 4 < 0$ induced paws which is impossible), and the total degree of G is $6n - 12$. Within these constraints, it is straightforward to see that $\sum_{v \in V(G)} d(v)^2$ is maximised if and only if two vertices of G have degree $n - 2$ and the rest have degree 4. Indeed, for any other degree sequence satisfying these conditions there exist $v_1, v_2 \in V(G)$ with $4 < d(v_1) < d(v_2) < n - 2$, so subtracting 1 from $d(v_1)$ and adding 1 to $d(v_2)$ gives a greater sum of squares. When this maximum occurs we have $\sum_{v \in V(G)} d(v)^2 = 2n^2 + 8n - 24$, so in fact this must be the degree sequence of G.

Let u and v be the two vertices in G of degree $n - 2$. The neighbourhood of u induces exactly a cycle, so $v \notin \Gamma(u)$ and hence $G = \overline{K_2} + C_{n-2}$. □

We are now ready to prove Theorem 5.1.

Proof of Theorem 5.1. We start by noting that the neighbourhood of every vertex in $\overline{K_2} + C_{n-2}$ induces a cycle, so for $n \geqslant 6$ the number of induced paws in the graph in which vertex v is the vertex of degree 3 is 0 if $d(v) = 4$ and $(n - 2)(n - 6)$ if $d(v) = n - 2$. It follows that there are exactly $2(n - 2)(n - 6) = 2n^2 - 16n + 24$ induced paws in $\overline{K_2} + C_{n-2}$ for all $n \geqslant 6$.

We will apply Lemma 5.11 with $s = 2$ to prove the theorem. It is clear that $\overline{K_2} + C_{n-2}$ is an n-vertex planar graph, so in the notation of the lemma we may take H to be the paw graph and $\mathcal{G}_n = \{\overline{K_2} + C_{n-2}\}$. Then $h(n) = 2n^2 - 16n + 24$ for $n \geqslant 6$ by the above, and for $n \geqslant 8$ we have $h(n) - h(n-1) = 4n - 18$ and $h(n) - h(n-2) = 8n - 40$. It follows from Lemmas 5.18 and 5.19 that condition $(\star)_2$ holds for these H and $(\mathcal{G}_n)$.

Therefore by Lemma 5.11, for sufficiently large n, $f_I(n, H) = 2n^2 - 16n + 24$ and $\overline{K_2} + C_{n-2}$ is the unique n-vertex planar graph which achieves this maximum. □

5.3 Induced copies of K_4^-

This section is devoted to proving Theorem 5.3, which concerns induced copies of K_4^-.

Proof of Theorem 5.3. We begin with the following claim.

Claim 5.21. *Let G be an n-vertex planar graph and suppose that $v \in V(G)$ has degree 3. Then v is in at most $n-2$ induced copies of K_4^-.*

Proof. Label the neighbours of v as u_1, u_2, and u_3. For $i \in [3]$, let

$$m_i = |(\Gamma(u_i) \cap \Gamma(u_{i+1})) \setminus \{v, u_{i+2}\}|,$$

where addition in the subscript is taken modulo 3. Note that, by the planarity of G, there is at most one vertex besides v in the common neighbourhood of u_1, u_2, and u_3, so $m_1 + m_2 \leqslant n-3$ and $m_1 + m_2 + m_3 \leqslant n-2$. Suppose that there exists an induced copy of K_4^- in G in which v is one of the vertices of degree 3. Then there is exactly one such induced copy of K_4^-, and without loss of generality $\Gamma(v)$ induces the path $u_1u_2u_3$. The number of induced copies of K_4^- in G in which v is one of the vertices of degree 2 is then at most $m_1 + m_2 \leqslant n-3$ as required. On the other hand, if there are no induced copies of K_4^- in G in which v is one of the vertices of degree 3, then there are at most $m_1 + m_2 + m_3 \leqslant n-2$ induced copies of K_4^- in G containing v. □

Using this analysis, we see that for $n \geqslant 5$ the vertices of degree 3 in $K_2 + P_{n-2}$ are each in exactly $n-2$ induced copies of K_4^-. Since $K_2 + P_2$ contains no induced copies of K_4^-, it follows that for $n \geqslant 4$ the number of induced copies of K_4^- in $K_2 + P_{n-2}$ is exactly

$$\sum_{i=5}^{n}(i-2) = \tfrac{1}{2}(n-4)(n+1) = \tfrac{1}{2}(n^2 - 3n - 4) =: g(n).$$

We now use the fact that $f_I(n, K_4^-) \geqslant g(n)$ to prove the following.

Claim 5.22. *For $n \geqslant 9$, every n-vertex planar graph containing $f_I(n, K_4^-)$ induced copies of K_4^- contains a separating triangle.*

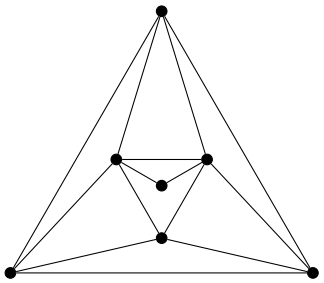


Figure 5.6: The only non-extremal n-vertex graph with more than $g(n)$ induced copies of K_4^-

Proof. Let $n \geqslant 9$ and let G be an n-vertex planar graph with no separating triangles. Since $n \geqslant 3$, the number of faces in any drawing of G, and hence the number of triangles in G, is at most $2n-4$. It is straightforward to see that any triangle in G is in at most three induced copies of K_4^-. The graph K_4^- contains two triangles, so G contains at most $3(2n-4)/2 = 3(n-2)$ induced copies of K_4^-, which is less than $g(n)$ for $n \geqslant 9$. The claim follows. □

Our proof will now proceed by induction on n. For $n \leqslant 10$ the theorem can be shown by a computer check. Moreover, a computer check[1] shows that the only non-extremal graph on $n \leqslant 9$ vertices which contains more than $g(n)$ induced copies of K_4^- is the 7-vertex graph shown in Figure 5.6, which contains 14 induced copies of K_4^-. This will be useful later.

Inductively, fix $n \geqslant 11$ and an extremal n-vertex planar graph G. It is straightforward to see that any vertex of degree at most 2 in G is in at most $n-3$ induced copies of K_4^-, so deleting such a vertex would yield an $(n-1)$-vertex planar graph with at least $f_I(n, K_4^-) - (n-3) \geqslant g(n) - n + 3 > g(n-1)$ induced copies of K_4^-. This contradicts the induction hypothesis, so $\delta(G) \geqslant 3$.

Suppose that $\delta(G) = 3$. By Claim 5.21, deleting a vertex of degree 3 from G yields an $(n-1)$-vertex planar graph with at least $f_I(n, K_4^-) - (n-2) \geqslant g(n) - n + 2 = g(n-1)$ induced copies of K_4^-. By the induction hypothesis, the only such graph is $K_2 + P_{n-3}$. It is easy to check that the only way to add a vertex of degree 3 to this graph so that this new vertex is in $n-2$ induced copies of K_4^- is to form $K_2 + P_{n-2}$. Hence, we may assume that $\delta(G) \geqslant 4$.

[1] As discussed in Section 5.1.3, our computer checks only run over connected planar graphs. However, for $n \leqslant 9$ there are clearly no disconnected n-vertex planar graphs containing more than $g(n)$ induced copies of K_4^-.

Let T be a separating triangle in G, which exists by Claim 5.22. Consider a drawing of G in which T has vertices in both its interior and exterior. Let k be the number of vertices in the interior of T in the drawing, so that $1 \leqslant k \leqslant n-4$. The number of induced copies of K_4^- using only vertices in T or its interior is at most $f_I(k+3, K_4^-)$ and the number using only vertices in T or its exterior is at most $f_I(n-k, K_4^-)$. One can show, using induction, that

$$\max_{1 \leqslant k \leqslant n-4} (f_I(k+3, K_4^-) + f_I(n-k, K_4^-)) < g(n) \tag{5.5}$$

for all $n \geqslant 11$. Indeed, note that inductively $f_I(k, K_4^-) \leqslant g(k)+3$ for all $4 \leqslant k \leqslant n-1$ except $k = 6$, for which we have $f_I(6, K_4^-) = g(6) + 5$. It follows that for $n \geqslant 11$, the left-hand side of (5.5) is at most the maximum of $g(k+3) + g(n-k) + 8$ over $1 \leqslant k \leqslant n-4$, which can be checked to be less than $g(n)$.

By (5.5), there must be at least one induced copy of K_4^- in G using vertices from both the interior and the exterior of T. We say that an induced copy of K_4^- *passes through* T if it has this property. Let $a_1 \geqslant a_2 \geqslant a_3$ be the sizes of the common neighbourhoods in the interior of T of each pair of vertices of T, and let b_1, b_2, and b_3 be the same for the exterior (where a_i and b_i relate to the same pair of vertices of T). Let G_1 be the subgraph of G induced on T and its interior, and let G_2 be that induced on T and its exterior. We now prove the following claim.

Claim 5.23. *The number of induced copies of K_4^- passing through T would strictly increase if G_2 were replaced by a (suitably positioned) copy of $K_2 + P_{n-k-2}$.*

Proof. Every induced copy of K_4^- passing through T contains exactly two vertices from T, and these are the vertices of degree 3 in the copy of K_4^-. It follows that the number of induced copies of K_4^- in G passing through T is $a_1b_1 + a_2b_2 + a_3b_3$. Since there is at least one induced copy of K_4^- in G passing through T, we deduce that $a_1 \geqslant 1$. If G_2 were replaced by a suitably positioned copy of $K_2 + P_{n-k-2}$, then the number of induced copies of K_4^- passing through T would be $a_1(n-k-3) + a_2 + a_3$. Thus, it is sufficient to show that $a_1b_1 + a_2b_2 + a_3b_3 < a_1(n-k-3) + a_2 + a_3$.

Suppose that there is no vertex in G_2 adjacent to all three vertices of T. Then $b_1 + b_2 + b_3 \leqslant n-k-3$. Moreover, since $\delta(G) \geqslant 4$, we have $b_1 < n-k-3$ by Lemma 5.12. Since $a_1 \geqslant a_2 \geqslant a_3$ and $a_1 \geqslant 1$, we therefore have

$$a_1b_1 + a_2b_2 + a_3b_3 \leqslant a_1(n-k-4) + a_2 < a_1(n-k-3) + a_2 + a_3,$$

as required.

Suppose instead that there does exist a vertex in G_2 adjacent to all three vertices of T. By the planarity of G there is exactly one such vertex. It is straightforward to use Lemma 5.12 and the fact that $\delta(G) \geqslant 4$ to show that $b_1 < n-k-3$, $b_1+b_2 < n-k-2$, and $b_1 + b_2 + b_3 < n - k - 1$. It follows that

$$a_1b_1 + a_2b_2 + a_3b_3 \leqslant a_1(n - k - 4) + a_2 + a_3 < a_1(n - k - 3) + a_2 + a_3,$$

where we have again used that $a_1 \geqslant a_2 \geqslant a_3$ and $a_1 \geqslant 1$. □

If G_2 contained at most $g(n-k)$ induced copies of K_4^-, then by Claim 5.23 we could replace it by a copy of $K_2 + P_{n-k-2}$ and increase the number of copies of K_4^- passing through T. Since $K_2 + P_{n-k-2}$ contains $g(n-k)$ induced copies of K_4^-, this means the total number of induced copies of K_4^- in the graph would increase, which is a contradiction. It follows (by symmetry) that G_1 and G_2 contain more than $g(k+3)$ and $g(n-k)$ induced copies of K_4^- respectively, so they are each one of the nine graphs shown in Figures 5.2 and 5.6.

The proof concludes by a computer check on these graphs: for each pair of graphs and each choice of a triangle that can bound a face in each, we identify these two triangles and count the number of induced copies of K_4^- in the resulting graph. No n-vertex graph so obtained contains at least $g(n)$ induced copies of K_4^-, which completes the proof of the theorem. □

5.4 Induced stars

In this section we prove the theorems concerning induced stars: Theorems 5.6, 5.7, 5.8, and 5.9.

5.4.1 Small stars

We begin with the proof of Theorem 5.6, which concerns the path on three vertices, P_3. We will need the following result which is well known and easily verifiable via Euler's formula.

Lemma 5.24. *For $n \geqslant 3$, every bipartite n-vertex planar graph has at most $2n-4$ edges.*

We now prove Theorem 5.6.

Proof of Theorem 5.6. It is easy to check that $K_{2,n-2}$ contains exactly $(n-2)^2$ induced copies of P_3, so $f_I(n, P_3) \geqslant (n-2)^2$ for all $n \geqslant 3$. The result is obvious for $n = 3$, so let $n \geqslant 4$ and let G be an n-vertex planar graph containing $f_I(n, P_3)$ induced copies of P_3. Inductively, every vertex of G is in at least $(n-2)^2 - (n-3)^2 = 2n-5$ induced copies of P_3, so $\delta(G) \geqslant 2$. We proceed by analysing how many times a vertex of given degree can appear as a leaf in an induced copy of P_3.

Claim 5.25. *Each vertex of degree $d \geqslant 2$ in G appears as a leaf in at most $2n-d-4$ induced copies of P_3.*

Proof. Let $v \in V(G)$ have degree $d \geqslant 2$. The number of induced copies of P_3 in which v appears as a leaf is the number of edges between the neighbourhood of v, $A \coloneqq \Gamma(v)$, and the second neighbourhood of v, $B \coloneqq \Gamma(\Gamma(v)) \setminus (\{v\} \cup \Gamma(v))$. The bipartite subgraph of G between A and $B \cup \{v\}$ contains at most $2n-4$ edges by Lemma 5.24, and since there are exactly d edges between v and A, this leaves at most $2n-d-4$ between A and B, as required. □

By the claim, the total number of induced copies of P_3 in G is at most $\frac{1}{2}\sum_v (2n - d(v) - 4) = n^2 - e(G) - 2n$. Since $f_I(n, P_3) \geqslant (n-2)^2 = n^2 - 4n + 4$, it follows that $e(G) \leqslant 2n-4$. The number of induced copies of P_3 containing a given edge uv of G is the number of vertices in $V(G) \setminus \{u, v\}$ with exactly one neighbour in $\{u, v\}$. This is clearly at most $n-2$. Since $f_I(n, P_3) \geqslant (n-2)^2$, it follows that $e(G) = 2n-4$, $f_I(n, P_3) = (n-2)^2$, and every edge of G is in exactly $n-2$ induced copies of P_3.

Let $v \in V(G)$ be of minimal degree, so that $d(v) \in \{2, 3\}$. Note that if $n \in \{4, 5\}$, then $d(v) = 2$. If $n \geqslant 6$ and $d(v) = 3$, then in order for the three edges incident to v to each appear in $n-2$ induced copies of P_3, the neighbours of v must all be adjacent to every vertex in $V(G) \setminus (\{v\} \cup \Gamma(v))$. This implies the existence of a subgraph of G isomorphic to $K_{3,3}$, which contradicts the planarity of G. Hence, for all $n \geqslant 4$, $d(v) = 2$. As above, the two neighbours of v must both be adjacent to every vertex in $V(G) \setminus (\{v\} \cup \Gamma(v))$, so G contains a $K_{2,n-2}$ subgraph. However, this accounts for all $2n-4$ edges of G, so in fact G is isomorphic to $K_{2,n-2}$, as required. □

The proof of Theorem 5.7, which concerns the three-leaf star, S_3, follows a similar approach.

Proof of Theorem 5.7. First, it is easy to check that all the graphs in the theorem statement contain the claimed number of induced copies of S_3. For $n \leqslant 8$ the result

can be shown by a computer check, so let $n \geqslant 9$. We have $f_I(n, S_3) \geqslant 2\binom{n-2}{3}$ for such n. Therefore, there is some vertex in G which appears as a leaf in at least

$$\frac{6}{n}\binom{n-2}{3} = n^2 - 9n + 26 - \frac{24}{n}$$

induced copies of S_3.

Clearly, any vertex of degree 1 appears in at most $\binom{n-2}{2} = \frac{1}{2}(n^2 - 5n + 6)$ induced copies of S_3, which is strictly less than $n^2 - 9n + 26 - \frac{24}{n}$ for $n \geqslant 9$. Hence, fix $v \in V(G)$ of degree $d \geqslant 2$ appearing as a leaf in at least $n^2 - 9n + 26 - \frac{24}{n}$ induced copies of S_3. Let $A \coloneqq \Gamma(v)$ and $B \coloneqq \Gamma(\Gamma(v)) \setminus (\{v\} \cup \Gamma(v))$ be the first and second neighbourhoods of v respectively, then for each $u \in A$ write $d'(u)$ for the number of edges between u and B. The number of induced copies of S_3 in which v appears as a leaf is then at most $\sum_{u \in A} \binom{d'(u)}{2}$. As in the proof of Theorem 5.6, the total number of edges between A and B is at most $2n - d - 4$, and clearly $d'(u) \leqslant n - d - 1$ for all $u \in A$. Under these conditions, the sum is maximised by assuming that two vertices of A have $n - d - 1$ edges to B, one has $d - 2$, and any others have none. This gives an upper bound on the number of times that v appears as a leaf in an induced copy of S_3 of

$$2\binom{n-d-1}{2} + \binom{d-2}{2} = \frac{3d^2}{2} - 2dn + \frac{d}{2} + n^2 - 3n + 5.$$

For fixed n, this expression is a quadratic in d with positive leading coefficient. At $d = 3$ it evaluates to $n^2 - 9n + 20$, and at $d = n - 1$ it evaluates to $\frac{1}{2}(n^2 - 7n + 12)$. Both of these expressions are strictly less than $n^2 - 9n + 26 - \frac{24}{n}$ for $n \geqslant 5$, from which we deduce that v has degree 2.

Let the neighbours of v be u and w. No induced copy of S_3 in G can contain all of u, v, and w, whence it is straightforward to see that v is in at most $2\binom{n-3}{2}$ induced copies of S_3. Inductively, every vertex of G is in at least $2\binom{n-2}{3} - 2\binom{n-3}{3} = 2\binom{n-3}{2}$ induced copies of S_3, so in fact v is in exactly this many. Applying the induction hypothesis again, we find that $G - v$ is isomorphic to $K_{2,n-3}$ or $K_2 + \overline{K_{n-3}}$. The only ways to add a vertex of degree 2 to these graphs while maintaining planarity and adding the required number of induced copies of S_3 is to form $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$ respectively. This completes the proof of the theorem. □

The proof of Theorem 5.8 is very similar to that of Theorem 5.7 and we essentially only include a sketch.

Proof of Theorem 5.8. The cases $n \leqslant 10$ follow from a computer check, then for $n \geqslant 11$ in any extremal graph there is a vertex appearing as a leaf in at least

$$\frac{8}{n}\binom{n-2}{4} = \frac{1}{3}\left(n^3 - 14n^2 + 71n - 154 + \frac{120}{n}\right)$$

induced copies of S_4. A vertex of degree 1 appears in at most $\binom{n-2}{3}$ induced copies of S_4, which is too few for $n \geqslant 11$. Similarly to above, a vertex of degree $d \geqslant 2$ appears as a leaf in at most

$$\begin{aligned} &2\binom{n-d-1}{3} + \binom{d-2}{3} \\ &= -\tfrac{1}{6}\big(d^3 - 6d^2n + 21d^2 + 6dn^2 - 24dn - 4d - 2n^3 + 12n^2 - 22n + 36\big) \end{aligned} \tag{5.6}$$

induced copies of S_4.

For fixed n, this expression is a cubic in d with a negative leading coefficient. Its partial derivative with respect to d is $-\frac{1}{6}(3d^2 - 12dn + 42d + 6n^2 - 24n - 4)$. At $d = 2n - 7$ this evaluates to $6n^2 - 60n + 151$, which is positive for all n. Since $2n - 7 \geqslant n - 1$ for $n \geqslant 6$, the expression in (5.6) is either monotonic in d between $d = 2$ and $d = n - 1$ or it is decreasing then increasing. At $d = 3$ it evaluates to $\frac{1}{3}(n^3 - 15n^2 + 74n - 120)$ and at $d = n - 1$ it evaluates to $\frac{1}{6}(n^3 - 12n^2 + 47n - 60)$. Both expressions are too small for $n \geqslant 6$. Hence, there exists a vertex of degree 2 in G. The proof now proceeds analogously to the S_3 case. □

5.4.2 Big stars

Our proof of Theorem 5.9 will be by an application of Lemma 5.11 for each $k \geqslant 3$, with $s = 1$ in each case. In the notation of that lemma we take $\mathcal{G}_n = \{K_{2,n-2}, K_2 + \overline{K_{n-2}}\}$, which gives $h(n) = 2\binom{n-2}{k}$, so that $h(n) - h(n-1) = 2\binom{n-3}{k-1}$. The following lemma is therefore appropriate for checking that condition $(\star)_1$ holds.

Lemma 5.26. *Let $k \geqslant 3$. Let n be large and suppose that G is an n-vertex planar graph in which every vertex is in at least $2\binom{n-3}{k-1}$ induced copies of S_k. Then G is $K_{2,n-2}$ or $K_2 + \overline{K_{n-2}}$.*

Once we have a lemma of this form it is straightforward to apply Lemma 5.11 to prove the theorem.

Proof of Theorem 5.9. We will apply Lemma 5.11 with $\mathcal{G}_n = \{K_{2,n-2}, K_2 + \overline{K_{n-2}}\}$ and $H = S_k$. Both graphs in $\mathcal{G}_n$ are planar and contain exactly $2\binom{n-2}{k}$ induced copies of S_k, so we may assign this value to $h(n)$. Condition $(\star)_1$ holds by Lemma 5.26, so

for large enough n, we have $f_I(n, S_k) = 2\binom{n-2}{k}$ and the only n-vertex planar graphs which achieve this maximum are $K_{2,n-2}$ and $K_2 + \overline{K_{n-2}}$. □

We now prove the lemma.

Proof of Lemma 5.26. By the planarity of G and the fact that $2\binom{n-3}{k-1} > 0$, we have $1 \leqslant \delta(G) \leqslant 5$. If $v \in V(G)$ has degree 1, then v is in at most $\binom{n-2}{k-1}$ induced copies of S_k, and $\binom{n-2}{k-1} < 2\binom{n-3}{k-1}$ if n is sufficiently large, so we have a contradiction. If $v \in V(G)$ has degree 2, then label its neighbours u and w and note that v is not the central vertex in any induced copy of S_k. If u and w are neighbours, then v is in at most $\binom{d(u)-2}{k-1} + \binom{d(w)-2}{k-1} \leqslant 2\binom{n-3}{k-1}$ induced copies of S_k. Equality holds, so $d(u) = d(w) = n-1$ and there are no edges in $G - \{u, w\}$, from which it follows that $G = K_2 + \overline{K_{n-2}}$. If instead u and w are not neighbours, then v is in at most $\binom{d(u)-1}{k-1} + \binom{d(w)-1}{k-1} \leqslant 2\binom{n-3}{k-1}$ induced copies of S_k. We have equality again, so $d(u) = d(w) = n-2$ and there are no edges in $G - \{u, w\}$. We deduce that $G = K_{2,n-2}$ in this case.

Now suppose $v \in V(G)$ has $d(v) = d \in \{3, 4, 5\}$. We will show that v is in fewer than $2\binom{n-3}{k-1}$ induced copies of S_k to obtain a contradiction. We may assume that every edge of G has exactly one endpoint in $\Gamma(v)$, since any edges with zero or two endpoints in $\Gamma(v)$ can be deleted without decreasing the number of induced copies of S_k containing v. If any vertex in $V(G) \setminus (\{v\} \cup \Gamma(v))$ is a neighbour of fewer than two neighbours of v, then this vertex has degree 0 or 1 in G, so we may add an edge from this vertex to a neighbour of v while preserving the planarity of G and not decreasing the number of induced copies of S_k containing v. We may therefore assume that every vertex in $V(G) \setminus (\{v\} \cup \Gamma(v))$ has at least two neighbours in $\Gamma(v)$. Label the neighbours of v as $u_1, \ldots, u_d$ where $d(u_1) \geqslant \ldots \geqslant d(u_d)$, then the number of induced copies of S_k in which v appears as a leaf is

$$\sum_{i=1}^{d} \binom{d(u_i) - 1}{k-1}.$$

Claim 5.27. *For $t \geqslant 2$ and $n_1 \geqslant n_2 \geqslant 1$,*

$$\binom{n_1 + 1}{t} + \binom{n_2 - 1}{t} \geqslant \binom{n_1}{t} + \binom{n_2}{t}.$$

Proof. We have $\binom{n_1+1}{t} = \binom{n_1}{t} + \binom{n_1}{t-1}$ and $\binom{n_2}{t} = \binom{n_2-1}{t} + \binom{n_2-1}{t-1}$, so the left-hand side of the inequality is

$$\binom{n_1}{t} + \binom{n_1}{t-1} + \binom{n_2-1}{t}$$

and the right-hand side is

$$\binom{n_1}{t} + \binom{n_2-1}{t} + \binom{n_2-1}{t-1}.$$

Thus the inequality holds true if and only if

$$\binom{n_1}{t-1} \geqslant \binom{n_2-1}{t-1},$$

so since $n_1 > n_2 - 1$, the claim is proved. □

Suppose that there exists $w \in V(G) \setminus (\{v\} \cup \Gamma(v))$ with exactly two neighbours in $\Gamma(v)$, say u_i and u_j for $i < j$, where $(i,j) \neq (1,2)$. Consider deleting w from G and replacing it with a vertex whose neighbours are exactly u_1 and u_2 (note that the resulting graph is not necessarily planar). In the case where $i = 1$, the change in the number of induced copies of S_k containing v is

$$\binom{d_G(u_2)}{k-1} - \binom{d_G(u_2)-1}{k-1} - \binom{d_G(u_j)-1}{k-1} + \binom{d_G(u_j)-2}{k-1}.$$

Since $d_G(u_2) \geqslant d_G(u_j)$, this is non-negative by the claim. Similarly, the number of induced copies of S_k containing v does not decrease in the cases $2 = i < j$ and $3 \leqslant i < j$.

Applying this repeatedly shows that v is in at most as many induced copies of S_k in G as it is in the graph constructed from G by replacing each vertex in $V(G)\setminus(\{v\}\cup\Gamma(v))$ that has exactly two neighbours in $\Gamma(v)$, with a vertex whose only neighbours are u_1 and u_2. Let this graph be F. Note that $d_F(u_1), d_F(u_2) \leqslant n - d \leqslant n - 3$ since neither vertex is adjacent to any of $u_1, \ldots, u_d$ in F. By the planarity of G, there exists an absolute constant c_0 such that there are at most c_0 vertices which have more than two common neighbours with v in G, so the same holds in F. Hence $d_F(u_3), \ldots, d_F(u_d) \leqslant c_0 + 1$. There also exists an absolute constant c_1 such that v occurs as the central vertex in at most c_1 induced copies of S_k in F. Thus the total number of induced copies of S_k in F which contain v is at most

$$c_1 + 3\binom{c_0}{k-1} + 2\binom{n-4}{k-1},$$

and this is also an upper bound for the number of such copies of S_k in G. Comparing the first and second coefficients of this expression considered as a polynomial in n of degree $k-1$, with the corresponding coefficients of $2\binom{n-3}{k-1}$ considered as another polynomial in n of degree $k-1$, we find that for large enough n there are fewer than $2\binom{n-3}{k-1}$ induced copies of S_k in G containing v, which gives the required contradiction and completes the proof of the lemma. □

5.5 Results on $f(n, H)$

In this section we prove our results in the not necessarily induced setting, namely Theorems 5.2 and 5.4.

5.5.1 Paws

We now give the proof of Theorem 5.2.

Proof of Theorem 5.2. We start by outlining one way of counting the number of paws in $K_2 + P_{n-2}$. First, note that $K_2 + P_2 = K_4$ contains 12 paw graphs. For $n \geqslant 5$, each vertex of degree 3 in $K_2 + P_{n-2}$ is in three paws in which it is the vertex of degree 3, $4n-8$ in which it is a vertex of degree 2, and $4n-11$ in which it is the vertex of degree 1. This makes a total of $8n-16$ paw graphs containing each vertex of degree 3 in $K_2 + P_{n-2}$. It follows that the total number of paws in $K_2 + P_{n-2}$ for $n \geqslant 5$ is

$$12 + \sum_{i=5}^{n}(8i-16) = 12 + \tfrac{1}{2}(n-4)(8n+8) = 4n^2 - 12n - 4.$$

One can check that the exceptional graphs A and B contain 108 and 156 paw graphs respectively, as required.

We proceed to prove the theorem by induction on n; the cases $n \in \{4, 5\}$ are easily verified. Fix $n \geqslant 6$ and let G be an n-vertex *maximal* planar graph containing $f(n, H)$ paws, where in this proof H denotes the paw graph. For each $v \in V(G)$, the number of paws in G in which v is the vertex of degree 3 is $(d(v)-2) \cdot e(G[\Gamma(v)])$. Since $G[\Gamma(v)]$ is an outerplanar graph, it has at most $2d(v)-3$ edges. The total number of paws in G is therefore at most

$$\begin{aligned}\sum_{v \in V(G)} (d(v)-2)(2d(v)-3) &= 2\sum_{v \in V(G)} d(v)^2 - 14 \cdot e(G) + 6n \\ &= 2\sum_{v \in V(G)} d(v)^2 - 36n + 84.\end{aligned}$$

If $\delta(G) \geqslant 4$, then we have $\sum_v d(v) = 6n - 12$ with $4 \leqslant d(v) \leqslant n - 1$ for all v. Under these conditions it is straightforward to see that $\sum_v d(v)^2$ is maximised by the degree sequence $(4, 4, 4, \ldots, 4, n - 3, n - 1)$. This translates to an upper bound of $4n^2 - 20n + 40$ on the number of paw graphs in G, which is less than $4n^2 - 12n - 4$ for $n \geqslant 6$. It follows that $\delta(G) = 3$.

Let $v \in V(G)$ have degree 3, and let $u_1, u_2, u_3 \in V(G)$ be its neighbours. We will upper bound the number of paws in G that contain v. For each $i \in [3]$, the number of paws containing the triangle vu_iu_{i+1} (where addition in the subscript is taken modulo 3) is $(d(v) - 2) + (d(u_i) - 2) + (d(u_{i+1}) - 2) = d(u_i) + d(u_{i+1}) - 3$. Hence, the number of paws in which v is not the pendant vertex is $2(d(u_1) + d(u_2) + d(u_3)) - 9$.

It was noted by Alon and Caro in [3] that the sum of the degrees of any three distinct vertices in an n-vertex planar graph is at most $2n + 2$. Indeed, since such a graph contains no $K_{3,3}$ subgraph, it contains at most two vertices adjacent to all three of the chosen vertices. The total degree of these chosen vertices is therefore at most $2 \cdot 3 + 2(n - 5) + 6 = 2n + 2$. It follows that the number of paws in G in which v is not the pendant vertex is at most $4n - 5$.

By Lemma 5.14, the number of *induced* paws in which v is the pendant vertex is at most $2n - 10$. The number of *non-induced* paws in which v is the pendant vertex and u_i is the vertex of degree 3 is at most

$$|(\Gamma(u_i) \cap \Gamma(u_{i+1})) \setminus \{v, u_{i+2}\}| + |(\Gamma(u_i) \cap \Gamma(u_{i+2})) \setminus \{v, u_{i+1}\}| + 1.$$

By the planarity of G, at most one vertex other than v is in the common neighbourhood of u_1, u_2, and u_3, from which it follows that the number of non-induced paws in which v is the pendant vertex is at most $2(n - 5) + 3 + 6 = 2n - 1$. Hence, the total number of paws in G containing v is at most $8n - 16$.

By the induction hypothesis and the count on the number of paws in $K_2 + P_{n-2}$, every vertex in G is in at least

$$(4n^2 - 12n - 4) - (4(n - 1)^2 - 12(n - 1) - 4) = 8n - 16$$

paws, so we have equality throughout the preceding paragraph. In particular, the neighbours of v in G have total degree $2n + 2$ in G, and hence total degree $2n - 1$ in $G - v$. Moreover, $G - v$ contains $f(n, H)$ paws.

Unless $n = 7$, the only way to add to $K_2 + P_{n-3}$ a vertex whose neighbours have total degree $2n - 1$ (in $K_2 + P_{n-3}$), while maintaining planarity, is to form $K_2 + P_{n-2}$.

If $n = 7$, then there are only two ways to achieve this: forming $K_2 + P_{n-2}$ or the exceptional graph A. The only way to add such a vertex to A is to form the exceptional graph B, and there is no way to add such a vertex to B. Therefore, by induction, the *maximal* planar graphs containing $f(n, H)$ paws are as claimed. To conclude, note that every extremal graph is a subgraph of an extremal graph which is also a maximal planar graph. However, every edge in every $K_2 + P_{n-2}$ and in both A and B is used in a paw graph, so there are no extremal graphs that are not maximal planar graphs. □

5.5.2 K_4^-

We will make use of the following result of Matolcsi and Nagy [104].

Lemma 5.28 ([104])**.** *For each $n \geqslant 3$, the maximum number of copies of P_3 in an n-vertex outerplanar graph is $\frac{1}{2}(n^2 + 3n - 12)$.*

We now give the proof of Theorem 5.4, which is similar to that of Theorem 5.2 given in the previous subsection.

Proof of Theorem 5.4. We start by briefly outlining one way to count the number of copies of K_4^- in $K_2 + P_{n-2}$. First, observe that $K_2 + P_2 = K_4$ contains exactly six copies of K_4^-. For $n \geqslant 5$, each vertex of degree 3 in $K_2 + P_{n-2}$ is in exactly three copies of K_4^- in which it is one of the vertices of degree 3, and exactly $n + 1$ in which it is one of the vertices of degree 2. Hence, the total number of copies of K_4^- in $K_2 + P_{n-2}$ is

$$6 + \sum_{i=5}^{n} (i + 4) = 6 + \tfrac{1}{2}(n - 4)(n + 13) = \tfrac{1}{2}(n^2 + 9n - 40).$$

One can also check that the exceptional graphs A and B contain exactly 36 and 48 copies of K_4^- respectively, as required.

We proceed by induction on n; the cases $n \in \{4, 5\}$ are straightforward to check. Fix $n \geqslant 6$ and let G be an n-vertex *maximal* planar graph containing $f(n, K_4^-)$ copies of K_4^-. For each $v \in V(G)$, the number of copies of K_4^- in G in which v is one of the vertices of degree 3 is the number of paths of length 2 in $G[\Gamma(v)]$. Since $G[\Gamma(v)]$ is outerplanar, the number of such paths is at most $\frac{1}{2}(d(v)^2 + 3d(v) - 12)$ by Lemma 5.28. Noting that K_4^- has two vertices of degree 3, it follows that the number of copies of

K_4^- in G is at most

$$\frac{1}{4}\sum_{v\in V(G)}(d(v)^2+3d(v)-12) = \frac{1}{4}\sum_{v\in V(G)} d(v)^2 + \frac{3}{2}\cdot e(G) - 3n$$
$$= \frac{1}{4}\sum_{v\in V(G)} d(v)^2 + \frac{3n}{2} - 9.$$

Suppose that $\delta(G) \geqslant 4$, then $\sum_v d(v) = 6n-12$ with $4 \leqslant d(v) \leqslant n-1$ for all v. It is straightforward to see that under these conditions $\sum_v d(v)^2$ is maximised by the degree sequence $(4,4,4,\ldots,4,n-3,n-1)$, which translates to an upper bound of $\frac{1}{2}(n^2+7n-29)$ on the number of copies of K_4^- in G. This is less than $\frac{1}{2}(n^2+9n-40)$ for $n \geqslant 6$, so in fact $\delta(G) = 3$.

Let $v \in V(G)$ have degree 3, and let u_1, u_2, and u_3 be its neighbours. The number of copies of K_4^- in G containing v is

$$3 + \sum_{i=1}^{3} |(\Gamma(u_i) \cap \Gamma(u_{i+1})) \setminus \{v\}|,$$

where addition in the subscript is taken modulo 3. By the planarity of G, there is at most one vertex other than v in the common neighbourhood of u_1, u_2, and u_3, so the number of copies of K_4^- in G containing v is at most $n+4$. By the induction hypothesis and the count on the number of copies of K_4^- in $K_2 + P_{n-2}$, every vertex in G is in at least

$$\tfrac{1}{2}(n^2+9n-40) - \tfrac{1}{2}((n-1)^2+9(n-1)-40) = n+4$$

copies of K_4^-, so in fact we have equality. In particular, every vertex of G is in the common neighbourhood of (at least) two neighbours of v. Moreover, $G-v$ contains $f(n,K_4^-)$ copies of K_4^-.

Unless $n=7$, the only way to add to K_2+P_{n-3} a vertex which has two common neighbours with every existing vertex, while maintaining planarity, is to form $K_2 + P_{n-2}$. If $n=7$, then this can be achieved only by forming K_2+P_{n-2} or the exceptional graph A. The only way to add such a vertex to A is to form the exceptional graph B, and there is no way to add such a vertex to B. Therefore, by induction, the *maximal* planar graphs containing $f(n,K_4^-)$ copies of K_4^- are as claimed. To conclude, note that every edge in these extremal graphs is contained in a copy of K_4^-, so there are no extremal graphs that are not maximal planar graphs. □

5.6 Some corollaries of the main theorems

In this section we prove some corollaries of Theorem 5.1 (the theorem on induced paws) and the theorems on induced stars. The proofs of these are all based on an idea applied by Alon and Caro [3], which is encapsulated in the following lemma.

Lemma 5.29. *Let H be a planar graph, and let F be an induced subgraph of H. Write f_1 for the number of induced copies of F in H. For each n, let $f_2(n)$ be the maximum possible number of induced copies of H in an n-vertex planar graph which contain a given copy of F as an induced subgraph. Then $f_I(n,H) \leqslant f_I(n,F)f_2(n)/f_1$. Moreover, if equality holds, then every n-vertex planar graph containing $f_I(n,H)$ induced copies of H contains $f_I(n,F)$ induced copies of F.*

Proof. Let G be a planar graph. Every induced copy of H in G contains exactly f_1 induced copies of F, and every induced copy of F in G is contained in at most $f_2(n)$ induced copies of H. Hence, the number of induced copies of H in G is at most $f_2(n)/f_1$ times the number of induced copies of F in G, from which the lemma follows. □

By Lemma 5.29, Theorem 5.1 has the following two corollaries.

Corollary 5.30. *Let H be the graph consisting of a 4-cycle with a fifth vertex joined to three of the vertices on the cycle. Then for large n we have $f_I(n,H) = n^2 - 8n + 12$ and the unique n-vertex planar graph which contains this many induced copies of H is $\overline{K_2} + C_{n-2}$.*

Proof. We will apply Lemma 5.29, where in the notation of the lemma we take F to be the paw graph and H to be as in the statement of the corollary. We then have $f_1 = 2$ and, since planar graphs contain no copies of $K_{3,3}$, $f_2(n) = 1$ (for $n \geqslant 5$). Now the lemma and Theorem 5.1 give $f_I(n,H) \leqslant n^2 - 8n + 12$ for large n. The corresponding lower bound is provided by $\overline{K_2} + C_{n-2}$, and by the 'moreover' part of the lemma this is the unique n-vertex planar graph achieving this maximum for large n. □

Corollary 5.31. *For $k \geqslant 2$, let H_k be the $(k+3)$-vertex graph consisting of a triangle, one of whose vertices is adjacent to k pendant vertices. Then*

$$f_I(n,H_k) = \frac{2n^{k+1}}{k!} + O(n^k).$$

Proof. For the lower bound, it is clear that $\overline{K_2} + C_{n-2}$ contains $2n^{k+1}/k! + O(n^k)$ induced copies of H_k. For the upper bound we apply Lemma 5.29 with H and F taken to be H_k and the paw graph respectively. This gives $f_1 = k$ and $f_2(n) = \binom{n-4}{k-1} = n^{k-1}/(k-1)! + O(n^{k-2})$, so by the lemma and Theorem 5.1, we have

$$f_I(n, H_k) \leqslant (2n^2 - 16n + 24) \cdot \left(\frac{n^{k-1}}{k!} + O(n^{k-2}) \right) = \frac{2n^{k+1}}{k!} + O(n^k),$$

as required. □

Next, we consider the complete bipartite graphs $K_{2,k}$ for $k \geqslant 2$. In the case $k = 2$, Hakimi and Schmeichel [70] showed that $f(n, C_4) = \frac{1}{2}(n^2 + 3n - 22)$ for all $n \geqslant 4$. Later, Alon and Caro [3] determined $f(n, H)$ in the cases where $H = K_{2,k}$ for $k \geqslant 3$ as follows.

Theorem 5.32 ([3])**.** *For $k \geqslant 3$ and $n \geqslant 5$ we have*

$$f(n, K_{2,k}) = \begin{cases} \binom{n-2}{k}, & \text{if } k \geqslant 5 \text{ or if } k = 4 \text{ and } n \neq 6 \\ \binom{n-2}{3} + 3(n-4), & \text{if } k = 3 \text{ and } n \neq 6 \\ 3, & \text{if } k = 4 \text{ and } n = 6 \\ 12, & \text{if } k = 3 \text{ and } n = 6. \end{cases}$$

In the induced setting, it was shown in [132] that $f_I(n, C_4) = \binom{n-2}{2}$ for sufficiently large n. Since $K_{2,n-2}$ contains $\binom{n-2}{k}$ induced copies of $K_{2,k}$ for each $k \geqslant 2$ and $n \geqslant 4$, it follows from Theorem 5.32 that $f_I(n, K_{2,k}) = \binom{n-2}{k}$ when $k \geqslant 5$ or $k = 4$ and $n \neq 6$. For $k = 4$ and $n = 6$, clearly $f_I(6, K_{2,4}) = 1$. In the case $k = 3$, it follows from Theorem 5.7 and Lemma 5.29 that $f_I(n, K_{2,3}) = \binom{n-2}{3}$ for $n \geqslant 7$. For $n = 5$ clearly $f_I(5, K_{2,3}) = 1$ and it is not much harder to check that for $n = 6$ we have $f_I(6, K_{2,3}) = 4$. In summary, we have the following.

Corollary 5.33. *For all $k \geqslant 3$ and $n \geqslant 5$ we have $f_I(n, K_{2,k}) = \binom{n-2}{k}$, and the same holds for sufficiently large n when $k = 2$.*

Another straightforward consequence of Theorem 5.9 and Lemma 5.29 is as follows.

Corollary 5.34. *Let $k \geqslant 3$. Then for large n we have*

$$f_I\left(n, K_2 + \overline{K_k}\right) = \binom{n-2}{k},$$

and the unique n-vertex planar graph containing this many induced copies of $K_2 + \overline{K_k}$ is $K_2 + \overline{K_{n-2}}$.

We can also use Lemma 5.29 to determine $f_I(n, H)$ asymptotically whenever H is a planar graph containing a common neighbourhood of size $|V(H)| - 2$, except in the cases where that common neighbourhood induces a path or a cycle.

Corollary 5.35. *Let $k \geqslant 2$ and let F be a proper spanning subgraph of P_k. Let ℓ be the number of different component sizes in F, let $a_1, a_2, \ldots, a_\ell$ be the number of components of each of the ℓ sizes, and let $m = \sum_{i=1}^{\ell} a_i$ be the total number of components of F. If H is either $K_2 + F$ or $\overline{K_2} + F$, then*

$$f_I(n, H) = \frac{n^m}{a_1!a_2! \ldots a_\ell!} + O(n^{m-1}).$$

Proof. First note that $K_2 + P_{n-2}$ contains $n^m/a_1!a_2! \ldots a_\ell! + O(n^{m-1})$ induced copies of K_2+F and $\overline{K_2}+C_{n-2}$ contains $n^m/a_1!a_2! \ldots a_\ell!+O(n^{m-1})$ induced copies of $\overline{K_2}+F$, so it is sufficient to prove the upper bounds. As in the statement of the corollary, let H be either K_2+F or $\overline{K_2}+F$. The proof will proceed by induction on k. If $k = 2$ then H is either C_4, in which case the claim follows from $f_I(n, C_4) = \frac{1}{2}(n^2 - 5n + 6)$ [132], or K_4^- in which case the claim follows from Theorem 5.3.

Now let $k \geqslant 3$ and assume the result holds for all smaller values of k. Let $k_1 \geqslant k_2 \geqslant \ldots \geqslant k_m$ be the sizes of the components of F. If $k_1 = 1$, then $H = K_{2,k}$ or $H = K_2+\overline{K_k}$ so the result follows from Corollary 5.33 or 5.34. Assuming therefore that $k_1 \geqslant 2$, let v be a vertex of degree 1 in a component of F of size k_1, and let $F' = F - v$. Then F' is a proper spanning subgraph of P_{k-1}. Let a_1 be the number of components of size k_1 in F, and let a' (possibly equal to 0) be the number of components of size $k_1 - 1$ in F. Let $H' = K_2 + F'$ if $H = K_2 + F$ and let $H' = \overline{K_2} + F'$ otherwise. Then by the induction hypothesis,

$$f_I(n, H') = \frac{n^m}{(a_1 - 1)!a_2! \ldots a_\ell!(a' + 1)} + O(n^{m-1}).$$

The graph H contains at least $2a_1$ induced copies of H' since deleting any vertex corresponding to a vertex of degree 1 in a component of F of size k_1 yields a copy of H'. Unless H' is a 4-cycle, the vertices of H' which correspond to the copy of K_2 or $\overline{K_2}$ are uniquely determined, so unless H is the graph considered in Corollary 5.30 (in which case the result follows immediately from that corollary), given an induced copy of H' in a planar graph G, there are at most $2(a' + 1)$ induced copies of H in G which contain the copy of H'. Indeed, the vertex we add to the copy of H' to obtain a copy of H must be adjacent to both of the vertices corresponding to the copy of K_2 or $\overline{K_2}$, and to a vertex corresponding to a vertex of degree 1 in a component of

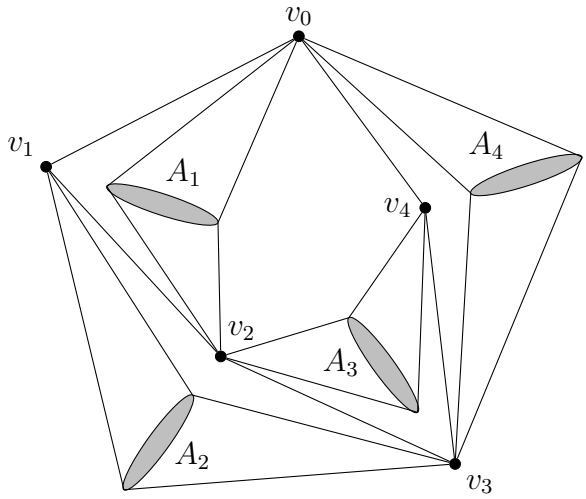


Figure 5.7: A planar graph containing many induced copies of P_4

F' of size $k_1 - 1$. Since G contains no subdivision of $K_{3,3}$, there are at most $2(a'+1)$ vertices of G satisfying these conditions. Hence, by Lemma 5.29,

$$f_I(n, H) \leqslant \frac{2(a'+1)}{2a_1} f_I(H', n) = \frac{n^m}{a_1! a_2! \dots a_\ell!} + O(n^{m-1}),$$

as required. □

5.7 Concluding remarks

The exact values of $f(n, H)$ and $f_I(n, H)$ are now known (at least for large n) for all connected graphs H on at most four vertices, except when H is the path on four vertices, P_4, in the induced version of the problem. It is not at all clear what the correct asymptotic value for $f_I(n, P_4)$ should be. The best construction we have found gives $f_I(n, P_4) \geqslant 9n^2/8 + O(n)$, but there is no particular reason to think that this is optimal.

The construction giving this lower bound is as follows. Start with a 5-cycle $v_0v_1v_2v_3v_4$ and partition the remaining vertices of the graph into four sets A_1, A_2, A_3, A_4 of as equal size as possible. For each $i \in [4]$, let the neighbours of each vertex in A_i be v_{i-1} and v_{i+1} (with subscripts taken modulo 5). See Figure 5.7 for an illustration of this graph. One can check that every vertex in each A_i is in $9n/4 + O(1)$ induced copies of P_4, so the total number of copies is $9n^2/8 + O(n)$. Note that it is possible to slightly increase the total number of induced copies of P_4 in this graph by, for example, adding a path through each of the sets A_i, but it is not clear how to improve on $9n^2/8 + O(n)$ as a lower bound.

We would be very interested to see a better construction or a proof that this one is asymptotically optimal.

Question 5.36. *What is the (asymptotic) value of $f_I(n, P_4)$ for large n?*

Bibliography

The numbers after the ↑ symbol indicate on which pages each item is referenced.

[1] M. Ajtai, J. Komlós, and E. Szemerédi. Largest random component of a k-cube. *Combinatorica*, **2**(1):1 – 7, 1982. https://doi.org/10.1007/BF02579276. ↑4

[2] A. F. Alameddine. On the number of cycles of length 4 in a maximal planar graph. *Journal of Graph Theory*, **4**(4):417 – 422, 1980. https://doi.org/10.1002/jgt.3190040411. ↑15

[3] N. Alon and Y. Caro. On the number of subgraphs of prescribed type of planar graphs with a given number of vertices. In M. Rosenfeld and J. Zaks, editors, *Annals of Discrete Mathematics (20): Convexity and Graph Theory*, volume 87 of *North-Holland Mathematics Studies*, pages 25 – 36. North-Holland, 1984. https://doi.org/10.1016/S0304-0208(08)72803-2. ↑14, 16, 116, 117, 118, 123, 146, 149, 150

[4] N. Alon, J. Pach, and J. Solymosi. Ramsey-type theorems with forbidden subgraphs. *Combinatorica*, **21**(2):155 – 170, 2001. https://doi.org/10.1007/s004930100016. ↑103

[5] N. Alon, K. Rudov, and L. Yariv. Dominance solvability in random games. *arXiv:2105.10743* preprint, 2021. ↑7

[6] B. Amiet, A. Collevecchio, and K. Hamza. When "better" is better than "best". *Operations Research Letters*, **49**(2):260 – 264, 2021. https://doi.org/10.1016/j.orl.2021.01.009. ↑85

[7] B. Amiet, A. Collevecchio, M. Scarsini, and Z. Zhong. Pure Nash equilibria and best-response dynamics in random games. *Mathematics of Operations Research*, **46**(4):1552 – 1572, 2021. https://doi.org/10.1287/moor.2020.1102. ↑7, 59, 86

[8] A. C. Antonir and A. Shapira. Bounding the number of odd paths in planar graphs via convex optimization. *Journal of Graph Theory*, **107**(2):330 – 343, 2024. https://doi.org/10.1002/jgt.23120. ↑16, 117

[9] K. R. Apt and S. Simon. A classification of weakly acyclic games. *Theory and Decision*, **78**(4):501 – 524, 2015. https://doi.org/10.1007/s11238-014-9436-1. ↑53

[10] I. Arieli and H. P. Young. Stochastic learning dynamics and speed of convergence in population games. *Econometrica*, **84**(2):627 – 676, 2016. https://doi.org/10.3982/ECTA10740. ↑86

[11] R. Arratia, L. Goldstein, and L. Gordon. Two moments suffice for Poisson approximations: the Chen-Stein method. *The Annals of Probability*, **17**(1):9 – 25, 1989. https://doi.org/10.1214/aop/1176991491. ↑54

[12] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 264 – 273, 2008. https://doi.org/10.1145/1386790.1386832. ↑86

[13] Y. Babichenko. Completely uncoupled dynamics and Nash equilibria. *Games and Economic Behavior*, **76**(1):1 – 14, 2012. https://doi.org/10.1016/j.geb.2012.06.004. ↑5, 8, 86

[14] Y. Babichenko. Best-reply dynamics in large binary-choice anonymous games. *Games and Economic Behavior*, **81**:130 – 144, 2013. https://doi.org/10.1016/j.geb.2013.04.007. ↑86

[15] R. C. Baker, G. Harman, and J. Pintz. The difference between consecutive primes, II. *Proceedings of the London Mathematical Society*, **83**(3):532 – 562, 2001. https://doi.org/10.1112/plms/83.3.532. ↑96

[16] P. Balister and B. Bollobás. Critical probabilities of 1-independent percolation models. *Combinatorics, Probability and Computing*, **21**(1-2):11 – 22, 2012. https://doi.org/10.1017/S0963548311000538. ↑2, 4, 18, 20, 46, 47, 49

[17] P. Balister, B. Bollobás, and M. Walters. Continuum percolation with steps in the square or the disc. *Random Structures & Algorithms*, **26**(4):392 – 403, 2005. https://doi.org/10.1002/rsa.20064. ↑2, 3, 19, 35, 37, 40, 41, 42, 43

[18] P. Balister, T. Johnston, M. Savery, and A. Scott. Improved bounds for 1-independent percolation on $\mathbb{Z}^n$. *arXiv:2206.12335* preprint, 2022. ↑iv, 3, 4, 5, 19, 20, 48, 49

[19] J. Balogh, P. Hu, B. Lidický, and F. Pfender. Maximum density of induced 5-cycle is achieved by an iterated blow-up of 5-cycle. *European Journal of Combinatorics*, **52**:47 – 58, 2016. https://doi.org/10.1016/j.ejc.2015.08.006. ↑14

[20] E. Berger, K. Choromanski, M. Chudnovsky, J. Fox, M. Loebl, A. Scott, P. Seymour, and S. Thomassé. Tournaments and colouring. *Journal of Combinatorial Theory, Series B*, **103**(1):1 – 20, 2013. https://doi.org/10.1016/j.jctb.2012.08.003. ↑103

[21] B. Bollobás. *Random Graphs*. Cambridge University Press, 2nd edition, 2001. https://doi.org/10.1017/CBO9780511814068. ↑57

[22] B. Bollobás, C. Gotsman, and E. Shamir. Connectivity and dynamics for random subgraphs of the directed cube. *Israel Journal of Mathematics*, **83**(3):321 – 328, 1993. https://doi.org/10.1007/BF02784059. ↑60, 63

[23] B. Bollobás, Y. Kohayakawa, and T. Łuczak. The evolution of random subgraphs of the cube. *Random Structures & Algorithms*, **3**(1):55 – 90, 1992. https://doi.org/10.1002/rsa.3240030106. ↑4

[24] B. Bollobás and O. Riordan. *Percolation*. Cambridge University Press, 2006. https://doi.org/10.1017/CBO9781139167383. ↑2

[25] C. Borgs, J. T. Chayes, R. van der Hofstad, G. Slade, and J. Spencer. Random subgraphs of finite graphs: I. The scaling window under the triangle condition. *Random Structures & Algorithms*, **27**(2):137 – 184, 2005. https://doi.org/10.1002/rsa.20051. ↑4

[26] C. Borgs, J. T. Chayes, R. van der Hofstad, G. Slade, and J. Spencer. Random subgraphs of finite graphs. II. The lace expansion and the triangle condition. *The Annals of Probability*, **33**(5):1886 – 1944, 2005. https://doi.org/10.1214/009117905000000260. ↑4

[27] C. Borgs, J. T. Chayes, R. van der Hofstad, G. Slade, and J. Spencer. Random subgraphs of finite graphs: III. The phase transition for the n-cube. *Combinatorica*, **26**(4):395 – 410, 2006. https://doi.org/10.1007/s00493-006-0022-1. ↑4

[28] R. C. Bose and S. Chowla. Theorems in the additive theory of numbers. *Commentarii Mathematici Helvetici*, **37**:141 – 147, 1962. https://doi.org/10.1007/BF02566968. ↑94, 96, 98

[29] A. T. Brandt. *Computational approaches in graph theory*. PhD thesis, University of Colorado Denver, 2016. Available at https://digital.auraria.edu/work/ns/35f95980-4e18-4613-b2a2-b3310f26cef1. ↑14

[30] M. Briański, J. Davies, and B. Walczak. Separating polynomial χ-boundedness from χ-boundedness. *arXiv:2201.08814v1* preprint, 2022. ↑11, 88, 91, 93

[31] M. Briański, J. Davies, and B. Walczak. Separating polynomial χ-boundedness from χ-boundedness. *Combinatorica*, **44**(1):1 – 8, 2024. https://doi.org/10.1007/s00493-023-00054-3. ↑12, 88, 93, 115

[32] R. L. Brooks. On colouring the nodes of a network. *Mathematical Proceedings of the Cambridge Philosophical Society*, **37**(2):194 – 197, 1941. https://doi.org/10.1017/S030500410002168X. ↑32

[33] A. Brouwer. Optimal packings of K_4's into a K_n. *Journal of Combinatorial Theory, Series A*, **26**(3):278 – 297, 1979. https://doi.org/10.1016/0097-3165(79)90105-5. ↑30

[34] O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo. Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, **36**(3):474 – 503, 2011. https://doi.org/10.1287/moor.1110.0500. ↑52

[35] A. Carbonero, P. Hompe, B. Moore, and S. Spirkl. A counterexample to a conjecture about triangle-free induced subgraphs of graphs with large chromatic number. *Journal of Combinatorial Theory, Series B*, **158**:63 – 69, 2023. https://doi.org/10.1016/j.jctb.2022.09.001. ↑11, 88, 91, 93, 115

[36] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006. https://doi.org/10.1017/CBO9780511546921. ↑8

[37] K. Coolsaet, S. D'hondt, and J. Goedgebeur. House of Graphs 2.0: A database of interesting graphs and more. *Discrete Applied Mathematics*, **325**:97 – 107, 2023. https://doi.org/10.1016/j.dam.2022.10.013. Available at https://houseofgraphs.org. ↑124

[38] C. Cox and R. R. Martin. Counting paths, cycles, and blow-ups in planar graphs. *Journal of Graph Theory*, **101**(3):521 – 558, 2022. https://doi.org/10.1002/jgt.22838. ↑15, 16, 116, 117

[39] C. Cox and R. R. Martin. The maximum number of 10- and 12-cycles in a planar graph. *Discrete Mathematics*, **346**(2):113245, 2023. https://doi.org/10.1016/j.disc.2022.113245. ↑15, 116

[40] A. N. Day, V. Falgas-Ravry, and R. Hancock. Long paths and connectivity in 1-independent random graphs. *Random Structures & Algorithms*, **57**(4):1007 – 1049, 2020. https://doi.org/10.1002/rsa.20972. ↑2, 3, 18, 45, 46, 48

[41] B. Descartes. k-chromatic graphs without triangles. *American Mathematical Monthly*, **61**(5):352 – 353, 1954. https://doi.org/10.1080/00029890.1954.11988473. ↑10

[42] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, **51**(1):161 – 166, 1950. https://doi.org/10.2307/1969503. ↑103

[43] M. Dresher. Probability of a pure equilibrium point in n-person games. *Journal of Combinatorial Theory*, **8**(1):134 – 145, 1970. https://doi.org/10.1016/S0021-9800(70)80015-1. ↑54

[44] G. Ellison. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, **67**(1):17 – 45, 2000. https://doi.org/10.1111/1467-937X.00119. ↑84

[45] D. Eppstein. Connectivity, graph minors, and subgraph multiplicity. *Journal of Graph Theory*, **17**(3):409 – 416, 1993. https://doi.org/10.1002/jgt.3190170314. ↑14

[46] P. Erdős and J. Spencer. Evolution of the n-cube. *Computers & Mathematics with Applications*, **5**(1):33 – 39, 1979. https://doi.org/10.1016/0898-1221(81)90137-1. ↑4

[47] P. Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, **11**:34 – 38, 1959. https://doi.org/10.4153/CJM-1959-003-9. ↑10, 92

[48] L. Esperet. Graph colorings, flows and perfect matchings. Habilitation thesis, Université Grenoble Alpes, 2017. Available at https://theses.hal.science/tel-01850463. ↑11

[49] C. Even-Zohar and N. Linial. A note on the inducibility of 4-vertex graphs. *Graphs and Combinatorics*, **31**:1367 – 1380, 2015. https://doi.org/10.1007/s00373-014-1475-4. ↑14

[50] A. Fabrikant, A. D. Jaggard, and M. Schapira. On the structure of weakly acyclic games. *Theory of Computing Systems*, **53**(1):107 – 122, 2013. https://doi.org/10.1007/s00224-013-9457-0. ↑52, 53

[51] V. Falgas-Ravry and V. Pfenninger. 1-independent percolation on $\mathbb{Z}^2 \times K_n$. *Random Structures & Algorithms*, **62**(4):887 – 910, 2023. https://doi.org/10.1002/rsa.21129. ↑2, 3, 4, 47, 49

[52] J. Folkman. Graphs with monochromatic complete subgraphs in every edge coloring. *SIAM Journal on Applied Mathematics*, **18**(1):19 – 24, 1970. https://doi.org/10.1137/0118004. ↑12, 13

[53] D. P. Foster and R. V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, **21**(1 – 2):40 – 55, 1997. https://doi.org/10.1006/game.1997.0595. ↑8

[54] D. P. Foster and H. P. Young. Regret testing: Learning to play Nash equilibrium without knowing you have an opponent. *Theoretical Economics*, **1**(3):341 – 367, 2006. Available at https://econtheory.org/ojs/index.php/te/article/viewArticle/20060341. ↑8, 10

[55] M. Freidlin, J. Szucs, and A. Wentzell. *Random Perturbations of Dynamical Systems*. Grundlehren der mathematischen Wissenschaften. Springer New York, 2012. https://doi.org/10.1007/978-3-642-25847-3. ↑84

[56] A. Frieze and M. Karoński. *Introduction to Random Graphs*. Cambridge University Press, 2015. https://doi.org/10.1017/CBO9781316339831. ↑62

[57] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, 1998. ↑5, 10

[58] D. Fudenberg and D. K. Levine. Conditional universal consistency. *Games and Economic Behavior*, **29**(1 – 2):104 – 130, 1999. https://doi.org/10.1006/game.1998.0705. ↑8

[59] F. Germano and G. Lugosi. Global Nash convergence of Foster and Young's regret testing. *Games and Economic Behavior*, **60**(1):135 – 154, 2007. https://doi.org/10.1016/j.geb.2006.06.001. ↑8

[60] D. Ghosh, E. Győri, R. R. Martin, A. Paulos, N. Salia, C. Xiao and O. Zamora. The maximum number of paths of length four in a planar graph. *Discrete Mathematics*, **344**(5):112317, 2021. https://doi.org/10.1016/j.disc.2021.112317. ↑15, 16, 117

[61] D. Ghosh, E. Győri, O. Janzer, A. Paulos, N. Salia, and O. Zamora. The maximum number of induced C_5's in a planar graph. *Journal of Graph Theory*, **99**(3):378 – 398, 2022. https://doi.org/10.1002/jgt.22745. ↑15, 116, 122

[62] A. Girão, F. Illingworth, E. Powierski, M. Savery, A. Scott, Y. Tamitegama, and J. Tan. Induced subgraphs of induced subgraphs of large chromatic number. *Combinatorica*, **44**(1):37 – 62, 2024. https://doi.org/10.1007/s00493-023-00061-4. ↑iv, 91

[63] K. Goldberg, A. Goldman, and M. Newman. The probability of an equilibrium point. *Journal of Research of the National Bureau of Standards*, **72**(2):93 – 101, 1968. https://doi.org/10.6028/jres.072b.012. ↑54

[64] A. Grzesik, E. Győri, A. Paulos, N. Salia, C. Tompkins, and O. Zamora. The maximum number of paths of length three in a planar graph. *Journal of Graph Theory*, **101**(3):493 – 510, 2022. https://doi.org/10.1002/jgt.22836. ↑16, 117

[65] A. Gyárfás. On Ramsey covering-numbers. In *Infinite and Finite Sets (Colloquium, Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), Vol. II*, Colloquia Mathematica Societatis János Bolyai, Vol. 10, pages 801 – 816. North-Holland, Amsterdam, 1975. ↑113

[66] A. Gyárfás. Problems from the world surrounding perfect graphs. *Applicationes Mathematicae*, **19**(3-4):413 – 441, 1987. https://doi.org/10.4064/am-19-3-4-413-441. ↑11, 113

[67] E. Győri, A. Paulos, N. Salia, C. Tompkins, and O. Zamora. Generalized planar Turán numbers. *The Electronic Journal of Combinatorics*, **28**:P4.32, 2021. https://doi.org/10.37236/9603. ↑15

[68] E. Győri, A. Paulos, N. Salia, C. Tompkins, and O. Zamora. The maximum number of pentagons in a planar graph. To appear in *Journal of Graph Theory*, 2024. https://doi.org/10.1002/jgt.23172. ↑15, 116

[69] A. Hajnal and P. Komjáth. Embedding graphs into colored graphs. *Transactions of the American Mathematical Society*, **307**(1):395 – 409, 1988. https://doi.org/10.2307/2000770. ↑13

[70] S. L. Hakimi and E. F. Schmeichel. On the number of cycles of length k in a maximal planar graph. *Journal of Graph Theory*, **3**(1):69 – 86, 1979. https://doi.org/10.1002/jgt.3190030108. ↑14, 116, 150

[71] T. E. Harris. A lower bound for the critical probability in a certain percolation process. *Mathematical Proceedings of the Cambridge Philosophical Society*, **56**(1):13 – 20, 1960. https://doi.org/10.1017/S0305004100034241. ↑4

[72] S. Hart. Adaptive heuristics. *Econometrica*, **73**(5):1401 – 1430, 2005. https://doi.org/10.1111/j.1468-0262.2005.00625.x. ↑10

[73] S. Hart and Y. Mansour. How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, **69**(1):107 – 126, 2010. https://doi.org/10.1016/j.geb.2007.12.002. ↑86

[74] S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, **68**(5):1127 – 1150, 2000. https://doi.org/10.1111/1468-0262.00153. ↑8, 10

[75] S. Hart and A. Mas-Colell. A reinforcement procedure leading to correlated equilibrium. In *Economics Essays*, pages 181 – 200. Springer, 2001. https://doi.org/10.1007/978-3-662-04623-4_12. ↑8

[76] S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, **93**(5):1830 – 1836, 2003. https://doi.org/10.1257/000282803322655581. ↑5, 8

[77] S. Hart and A. Mas-Colell. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, **57**(2):286 – 303, 2006. https://doi.org/10.1016/j.geb.2005.09.007. ↑5, 8, 80, 81

[78] S. Hart and A. Mas-Colell. *Simple Adaptive Strategies: From Regret-Matching to Uncoupled Dynamics*. World Scientific, 2013. https://doi.org/10.1142/8408. ↑5

[79] H. Hatami, J. Hirst, and S. Norine. The inducibility of blow-up graphs. *Journal of Combinatorial Theory, Series B*, **109**:196 – 212, 2014. https://doi.org/10.1016/j.jctb.2014.06.005. ↑14

[80] E. Heath, R. R. Martin, and C. Wells. The maximum number of odd cycles in a planar graph. *arXiv:2307.00116* preprint, 2023. ↑15, 116

[81] T. Heinrich, Y. Jang, L. Mungo, M. Pangallo, A. Scott, B. Tarbush, and S. Wiese. Best-response dynamics, playing sequences, and convergence to equilibrium in random games. *International Journal of Game Theory*, **52**:703 – 735, 2023. https://doi.org/10.1007/s00182-023-00837-4. ↑83, 85

[82] M. Heydenreich and R. van der Hofstad. *Progress in High-Dimensional Percolation and Random Graphs*. Springer, 2017. https://doi.org/10.1007/978-3-319-62473-0. ↑4, 5, 48

[83] J. Hofbauer and W. H. Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, **70**(6):2265 – 2294, 2002. https://doi.org/10.1111/j.1468-0262.2002.00440.x. ↑7, 53

[84] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998. https://doi.org/10.1017/CBO9781139173179. ↑10

[85] T. Huynh, G. Joret, and D. R. Wood. Subgraph densities in a surface. *Combinatorics, Probability and Computing*, **31**(5):812 – 839, 2022. https://doi.org/10.1017/S0963548321000560. ↑14, 116

[86] A. D. Jaggard, N. Lutz, M. Schapira, and R. N. Wright. Self-stabilizing uncoupled dynamics. In *Algorithmic Game Theory: 7th International Symposium, SAGT 2014, Proceedings 7*, pages 74 – 85. Springer, 2014. https://doi.org/10.1007/978-3-662-44803-8_7. ↑5, 8, 80, 81

[87] T. Johnston, M. Savery, A. Scott, and B. Tarbush. Game connectivity and adaptive dynamics. *arXiv:2309.10609* preprint, 2023. ↑iv, 10, 54

[88] E. Kalai and E. Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, **61**(5):1019 – 1045, 1993. https://doi.org/10.2307/2951492. ↑10

[89] M. Kandori, G. J. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, **61**(1):29 – 56, 1993. https://doi.org/10.2307/2951777. ↑84

[90] M. Kandori and R. Rob. Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory*, **65**(2):383 – 414, 1995. https://doi.org/10.1006/jeth.1995.1014. ↑84

[91] H. Kesten. The critical probability of bond percolation on the square lattice equals 1/2. *Communications in Mathematical Physics*, **74**:41 – 59, 1980. https://doi.org/10.1007/BF01197577. ↑4

[92] H. A. Kierstead. Classes of graphs that are not vertex Ramsey. *SIAM Journal on Discrete Mathematics*, **10**(3):373 – 380, 1997. https://doi.org/10.1137/S0895480194264769. ↑13

[93] H. A. Kierstead and W. T. Trotter. Colorful induced subgraphs. *Discrete Mathematics*, **101**:165 – 169, 1992. https://doi.org/10.1016/0012-365X(92)90600-K. ↑11, 91, 93

[94] H. A. Kierstead and Y. Zhu. Classes of graphs that exclude a tree and a clique and are not vertex Ramsey. *Combinatorica*, **16**(4):493 – 504, 1996. https://doi.org/10.1007/BF01271268. ↑13

[95] P. Komjáth and V. Rödl. Coloring of universal graphs. *Graphs and Combinatorics*, **2**(1):55 – 60, 1986. https://doi.org/10.1007/BF01788077. ↑13

[96] P. Komjáth and S. Shelah. A consistent edge partition theorem for infinite graphs. *Acta Mathematica Hungarica*, **61**(1-2):115 – 120, 1993. https://doi.org/10.1007/BF01872104. ↑13

[97] D. Král', S. Norin, and J. Volec. A bound on the inducibility of cycles. *Journal of Combinatorial Theory, Series A*, **161**:359 – 363, 2019. https://doi.org/10.1016/j.jcta.2018.08.003. ↑14

[98] C. Kuratowski. Sur le problème des courbes gauches en topologie. *Fundamenta Mathematicae*, **15**:271 – 283, 1930. https://doi.org/10.4064/fm-15-1-271-283. ↑123, 124

[99] T. M. Liggett, R. H. Schonmann, and A. M. Stacey. Domination by product measures. *The Annals of Probability*, **25**(1):71 – 95, 1997. https://doi.org/10.1214/aop/1024404279. ↑2, 35

[100] C.-H. Liu. Homomorphism counts in robustly sparse graphs. *arXiv:2107.00874* preprint, 2021. ↑14, 116

[101] Z. Lv, E. Győri, Z. He, N. Salia, C. Tompkins, and X. Zhu. The maximum number of copies of an even cycle in a planar graph. *Journal of Combinatorial Theory, Series B*, **167**:15 – 22, 2024. https://doi.org/10.1016/j.jctb.2024.01.005. ↑15, 116

[102] J. R. Marden, G. Arslan, and J. S. Shamma. Regret based dynamics: convergence in weakly acyclic games. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1 – 8, 2007. https://doi.org/10.1145/1329125.1329175. ↑83

[103] J. R. Marden, H. P. Young, G. Arslan, and J. S. Shamma. Payoff-based dynamics for multiplayer weakly acyclic games. *SIAM Journal on Control and Optimization*, **48**(1):373 – 396, 2009. https://doi.org/10.1137/070680199. ↑83

[104] D. Matolcsi and Z. L. Nagy. Generalised outerplanar Turán numbers and maximum number of k-vertex subtrees. *Discrete Applied Mathematics*, **307**:115 – 124, 2022. https://doi.org/10.1016/j.dam.2021.10.017. ↑147

[105] C. McDiarmid, A. Scott, and P. Withers. The component structure of dense random subgraphs of the hypercube. *Random Structures & Algorithms*, **59**(1):3 – 24, 2021. https://doi.org/10.1002/rsa.20990. ↑5, 7, 26, 59, 60, 63, 66, 67

[106] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. https://doi.org/10.1017/CBO9780511813603. ↑21

[107] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, **14**(1):124 – 143, 1996. https://doi.org/10.1006/game.1996.0044. ↑7, 53

[108] J. Nešetřil. *Teorie Grafů*. Praha: Státní Nakladatelství Technické Literatury, 1st edition, 1979. ↑11, 88

[109] J. Nešetřil. Ramsey theory. In *Handbook of Combinatorics*, volume 2, pages 1331 – 1403. MIT Press, 1996. https://doi.org/10.5555/233228.233234. ↑12

[110] J. Nešetřil, 2023. Personal communication. ↑106, 115

[111] J. Nešetřil and P. Ossona de Mendez. *Sparsity*. Springer Berlin, Heidelberg, 2012. https://doi.org/10.1007/978-3-642-27875-4. ↑92

[112] J. Nešetřil and V. Rödl. Partitions of vertices. *Commentationes Mathematicae Universitatis Carolinae*, **17**(1):85 – 95, 1976. Available at http://dml.cz/dmlcz/105676. ↑12, 114

[113] J. Nešetřil and V. Rödl. The Ramsey property for graphs with forbidden complete subgraphs. *Journal of Combinatorial Theory, Series B*, **20**(3):243 – 249, 1976. https://doi.org/10.1016/0095-8956(76)90015-0. ↑13

[114] J. Nešetřil and V. Rödl. A short proof of the existence of highly chromatic hypergraphs without short cycles. *Journal of Combinatorial Theory, Series B*, **27**(2):225 – 227, 1979. https://doi.org/10.1016/0095-8956(79)90084-4. ↑92, 106

[115] J. Newton and R. Sawa. Conventions and social choice in large games, 2024. Preprint available at SSRN: https://doi.org/10.2139/ssrn.4775575. ↑9, 53, 84

[116] N. Pippinger and M. C. Golumbic. The inducibility of graphs. *Journal of Combinatorial Theory, Series B*, **19**(3):189 – 203, 1975. https://doi.org/10.1016/0095-8956(75)90084-2. ↑13

[117] I. Y. Powers. Limiting distributions of the number of pure strategy Nash equilibria in n-person games. *International Journal of Game Theory*, **19**(3):277 – 286, 1990. https://doi.org/10.1007/BF01755478. ↑54

[118] B. S. Pradelski and H. P. Young. Learning efficient Nash equilibria in distributed systems. *Games and Economic Behavior*, **75**(2):882 – 897, 2012. https://doi.org/10.1016/j.geb.2012.02.017. ↑8, 86

[119] Y. Rinott and M. Scarsini. On the number of pure strategy Nash equilibria in random games. *Games and Economic Behavior*, **33**(2):274 – 293, 2000. https://doi.org/10.1006/game.1999.0775. ↑54, 57, 75

[120] V. Rödl. On the chromatic number of subgraphs of a given graph. *Proceedings of the American Mathematical Society*, **64**:370 – 371, 1977. https://doi.org/10.1090/S0002-9939-1977-0469806-4. ↑13

[121] V. Rödl and N. Sauer. The Ramsey property for families of graphs which exclude a given graph. *Canadian Journal of Mathematics*, **44**(5):1050 – 1060, 1992. https://doi.org/10.4153/CJM-1992-064-7. ↑12

[122] V. Rödl, N. Sauer, and X. Zhu. Ramsey families which exclude a graph. *Combinatorica*, **15**(4):589 – 596, 1995. https://doi.org/10.1007/BF01192529. ↑12

[123] V. Rödl and P. Winkler. A Ramsey-type theorem for orderings of a graph. *SIAM Journal on Discrete Mathematics*, **2**(3):402 – 406, 1989. https://doi.org/10.1137/0402035. ↑103

[124] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, **2**(1):65 – 67, 1973. https://doi.org/10.1007/BF01737559. ↑7

[125] T. Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016. https://doi.org/10.1017/CBO9781316779309. ↑7, 53, 80

[126] T. Roughgarden. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2021. https://doi.org/10.1017/9781108637435. ↑9

[127] K. Rudov. Fragile stable matchings. *arXiv:2403.12183* preprint, 2024. ↑84

[128] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010. ↑10

[129] N. Sauer. On the Ramsey property of families of graphs. *Transactions of the Amertican Mathematical Society*, **374**(3):785 – 833, 1995. https://doi.org/10.2307/2154874. ↑12

[130] N. W. Sauer and X. Zhu. Graphs which do not embed a given graph and the Ramsey property. In *Sets, Graphs and Numbers (Budapest, 1991)*, volume 60 of *Colloquia Mathematica Societatis János Bolyai*, pages 631 – 636. North-Holland, Amsterdam, 1992. ↑12

[131] M. Savery. Planar graphs with the maximum number of induced 6-cycles. *The Electronic Journal of Combinatorics*, **30**:P4.6, 2023. https://doi.org/10.37236/11944. ↑15, 116, 122

[132] M. Savery. Planar graphs with the maximum number of induced 4-cycles or 5-cycles. *Graphs and Combinatorics*, **40**(3):46, 2024. https://doi.org/10.1007/s00373-024-02773-w. ↑15, 116, 122, 150, 151

[133] A. Scott. Graphs of large chromatic number. In *Proceedings of the International Congress of Mathematicians 2022*, volume 6, pages 4660 – 4680, 2023. https://doi.org/10.4171/ICM2022/149. ↑11

[134] A. Scott and P. Seymour. Induced subgraphs of graphs with large chromatic number. I. Odd holes. *Journal of Combinatorial Theory, Series B*, **121**:68 – 84, 2016. https://doi.org/10.1016/j.jctb.2015.10.002. ↑115

[135] A. Scott and P. Seymour. Induced subgraphs of graphs with large chromatic number. VI. Banana trees. *Journal of Combinatorial Theory, Series B*, **145**:487 – 510, 2020. https://doi.org/10.1016/j.jctb.2020.01.004. ↑12

[136] A. Scott and P. Seymour. A survey of χ-boundedness. *Journal of Graph Theory*, **95**(3):473 – 504, 2020. https://doi.org/10.1002/jgt.22601. ↑11, 88

[137] S. Sidon. Ein Satz über trigonometrische Polynome und seine Anwendung in der Theorie der Fourier-Reihen. *Mathematische Annalen*, **106**:536 – 539, 1932. https://doi.org/10.1007/BF01455900. ↑94

[138] W. Stanford. A note on the probability of k pure Nash equilibria in matrix games. *Games and Economic Behavior*, **9**(2):238 – 246, 1995. https://doi.org/10.1006/game.1995.1019. ↑54

[139] D. P. Sumner. Subtrees of a graph and the chromatic number. In *The Theory and Applications of Graphs (Kalamazoo, Michigan, 1980)*, pages 557 – 576. Wiley, New York, 1981. ↑113

[140] B. Swenson, C. Eksin, S. Kar, and A. Ribeiro. Distributed inertial best-response dynamics. *IEEE Transactions on Automatic Control*, **63**(12):4294 – 4300, 2018. https://doi.org/10.1109/TAC.2018.2817161. ↑82

[141] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arithmetica*, **27**:199 – 245, 1975. https://doi.org/10.4064/aa-27-1-199-245. ↑13

[142] S. Takahashi and T. Yamamori. The pure Nash equilibrium property and the quasi-acyclic condition. *Economics Bulletin*, **3**(22):1 – 6, 2002. Available at http://www.accessecon.com/pubs/EB/2002/Volume3/EB-02C70011A.pdf. ↑53

[143] J. van den Berg and A. Ermakov. A new lower bound for the critical probability of site percolation on the square lattice. *Random Structures & Algorithms*, **8**(3):199 – 212, 1996. https://doi.org/10.1002/(SICI)1098-2418(199605)8:3<199::AID-RSA4>3.0.CO;2-T. ↑3, 48

[144] R. van der Hofstad and A. Nachmias. Unlacing hypercube percolation: a survey. *Metrika*, **77**(1):23 – 50, 2014. https://doi.org/10.1007/s00184-013-0473-5. ↑4

[145] R. van der Hofstad and G. Slade. Asymptotic expansions in n^{-1} for percolation critical values on the n-cube and $\mathbb{Z}^n$. *Random Structures & Algorithms*, **27**(3):331 – 357, 2005. https://doi.org/10.1002/rsa.20074. ↑4

[146] R. van der Hofstad and G. Slade. Expansion in n^{-1} for percolation critical values on the n-cube and $\mathbb{Z}^n$: the first three terms. *Combinatorics, Probability and Computing*, **15**(5):695 – 713, 2006. https://doi.org/10.1017/S0963548306007498. ↑4

[147] B. L. van der Waerden. Beweis einer Baudetschen Vermutung. *Nieuw Archief voor Wiskunde*, **15**:212 – 216, 1927. Available at https://www.delpher.nl/nl/tijdschriften/view?coll=dts&identifier=MMKWG01%3A022157001%3A00228. ↑13

[148] J. W. Weibull. *Evolutionary Game Theory*. MIT Press, 1997. ↑10

[149] R. M. Wilson. The construction of group divisible designs and partial planes having the maximum number of lines of a given size. In *Proceedings of The Second Chapel Hill Conference on Combinatorial Mathematics and its Applications, University of North Carolina, Chapel Hill, NC*, pages 488 – 497, 1970. ↑30

[150] D. R. Wood. On the maximum number of cliques in a graph. *Graphs and Combinatorics*, **23**:337 – 352, 2007. https://doi.org/10.1007/s00373-007-0738-8. ↑14

[151] N. C. Wormald. On the frequency of 3-connected subgraphs of planar graphs. *Bulletin of the Australian Mathematical Society*, **34**(2):309 – 317, 1986. https://doi.org/10.1017/S0004972700010182. ↑14

[152] H. P. Young. The evolution of conventions. *Econometrica*, **61**(1):57 – 84, 1993. https://doi.org/10.2307/2951778. ↑6, 10, 53, 83, 84

[153] H. P. Young. *Strategic Learning and its Limits*. Oxford University Press, 2004. https://doi.org/10.1093/acprof:oso/9780199269181.001.0001. ↑5, 9, 81, 82

[154] H. P. Young. The possible and the impossible in multi-agent learning. *Artificial Intelligence*, **171**(7):429 – 433, 2007. https://doi.org/10.1016/j.artint.2006.10.015. ↑5, 8

[155] H. P. Young. Learning by trial and error. *Games and Economic Behavior*, **65**(2):626 – 643, 2009. https://doi.org/10.1016/j.geb.2008.02.011. ↑8, 10, 80, 81, 82

[156] R. Yuster. On the exact maximum induced density of almost all graphs and their inducibility. *Journal of Combinatorial Theory, Series B*, **136**:81 – 109, 2019. https://doi.org/10.1016/j.jctb.2018.09.005. ↑14

[157] R. M. Ziff. Spanning probability in 2D percolation. *Physical Review Letters*, **69**(18):2670 – 2673, 1992. https://doi.org/10.1103/PhysRevLett.69.2670. ↑3, 48

[158] A. Zykov. On some properties of linear complexes (in Russian). *Matematicheskiĭ Sbornik*, **24**:163 – 188, 1949. For an English translation see *American Mathematical Society Translations*, **79**, 1952. ↑10, 91