Forum of Mathematics, Sigma (2026), **Vol. 14**:e32 1–57
doi:10.1017/fms.2026.10178

CAMBRIDGE UNIVERSITY PRESS

RESEARCH ARTICLE

Sharp thresholds for Ramsey properties

Ehud Friedgut[1], Eden Kuperwasser[2], Wojciech Samotij[3] and Mathias Schacht[4]

[1]Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel;
E-mail: ehud.friedgut@weizmann.ac.il.
[2]School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel;
E-mail: kuperwasser@mail.tau.ac.il (Corresponding author).
[3]School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel; E-mail: samotij@tauex.tau.ac.il.
[4]Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany; E-mail: schacht@math.uni-hamburg.de.

Received: 15 August 2022; **Revised:** 28 August 2025; **Accepted:** 28 August 2025

2020 Mathematics Subject Classification: *Primary* – 05C80, 05D10

Abstract

In this work, we develop a unified framework for establishing sharp threshold results for various Ramsey properties. To achieve this, we view such properties as noncolourability of auxiliary hypergraphs. Our main technical result gives sufficient conditions on a sequence of such hypergraphs that guarantee that this noncolourability property has a sharp threshold in subhypergraphs induced by random subsets of the vertices.

Furthermore, we verify these conditions in several cases of interest. In the classical setting of Ramsey theory for graphs, we show that the property of being Ramsey for a graph H in r colours has a sharp threshold in $G_{n,p}$, for all $r \geqslant 2$ and all H in a class of graphs that includes all cliques and cycles. In the arithmetic setting, we establish sharpness of thresholds for the properties corresponding to van der Waerden's theorem and Schur's theorem, also in any number of colours.

Contents

1 Introduction 1
2 The main result 6
3 An outline of the proof 10
4 Preliminaries and tools 15
5 Containers for colourings 19
6 Rainbow stars and constellations 30
7 Applications 41
A Weak threshold – the 1-statement 53
B A container lemma for sparse sets 55
C A graph without the star-constellation property 56
References 57

1. Introduction

A typical result in Ramsey theory states that, given a structure A and an integer $r \geqslant 2$, every colouring of the elements of any sufficiently 'rich' set V with r colours must contain a monochromatic copy of A. The

This research was supported by the grant I-1358-304.6/2016 from the German–Israeli Foundation for Scientific Research and Development (GIF). Additionally, EF was supported by the Minerva Stiftung with funding from the Federal German Ministry for Education and Research, and by the Israeli Science Foundation Grant number 2364/25.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

most prominent example is Ramsey's theorem [22], which states that, for any graph H and any integer $r \geqslant 2$, every r-colouring of the edges of a sufficiently large complete graph must yield a monochromatic copy of H. Two other famous instances, which actually predate [22], include van der Waerden's theorem [31] on arithmetic progressions and Schur's theorem [30] on additive triples.

In the 1980s, researchers have turned to studying Ramsey properties of random sets while trying to better understand what 'richness' assumptions a set V needs to satisfy so that it contains a monochromatic copy of a given structure A in every r-colouring. The seminal work of Frankl and Rödl [7] proves the existence of a K_4-free graph whose every 2-colouring contains a monochromatic triangle by considering the binomial random graph $G_{n,p}$ for an appropriately chosen edge density p. Soon afterwards, Łuczak, Ruciński, and Voigt [19] initiated the systematic study of Ramsey properties of random graphs, which has quickly become one of the central topics in probabilistic combinatorics.

Given a finite set V and a $p \in [0,1]$, we will write V_p to denote the random subset of V obtained by independently retaining each of its elements with probability p. While investigating, for a sequence of sets V whose sizes grow to infinity, the probability that V_p has a given property $\mathcal{P}$, one naturally encounters threshold phenomena. We say that a sequence of probabilities $\hat{p}$ is a *threshold* for a property $\mathcal{P}$ if the following two statements hold: On the one hand, for any $p \ll \hat{p}$, the probability that $V_p \in \mathcal{P}$ tends to zero; on the other hand, for any $p \gg \hat{p}$, the same probability tends to one. These two statements are aptly termed the 0-statement and the 1-statement, respectively. Thresholds have been a central theme in probabilistic combinatorics since its very inception and date back to the seminal paper of Erdős and Rényi [5] which initiated the systematic study of random graphs.

The celebrated theorem of Bollobás and Thomason [4] asserts the existence of a threshold for any property of sets that is monotone and nontrivial; this includes all Ramsey properties. However, this general theorem provides little clue regarding the location of this threshold. As a result, the main focus of the vast majority of the many works on Ramsey properties of random sets was locating the corresponding threshold. In particular, the locations of the thresholds for all of the aforementioned Ramsey properties were discovered in a series of papers by Graham, Rödl, and Ruciński [14], Rödl and Ruciński [23, 24, 25, 26], and Friedgut, Rödl, and Schacht [13]. Actually, these papers went one step further and showed that the 0-statement and the 1-statement hold already when $p \leqslant c_0 \cdot \hat{p}$ and $p \geqslant c_1 \cdot \hat{p}$, respectively, for some sequence $\hat{p}$ and positive constants c_0 and c_1.

It is very natural to ask whether this gap can be reduced even further. A property is said to have a *sharp threshold* if, for some threshold $\hat{p}$ and every positive ε, the 0-statement holds for $p \leqslant (1-\varepsilon)\hat{p}$ whereas the 1-statement holds for $p \geqslant (1+\varepsilon)\hat{p}$; otherwise, we say that the property has a *coarse threshold*. The notion of sharpness is closely reminiscent of the physical phenomenon of phase transition, where certain types of matter undergo a profound change in behaviour when their temperature crosses a certain point. Sharpness of thresholds has been established for several natural graph properties, such as connectivity, the existence of a perfect matching, and Hamiltonicity. On the other hand, many properties have been shown to have only a coarse threshold.

The presence of a sharp threshold, or a lack thereof, was demystified in the work of Friedgut [8]. Roughly speaking, the main result of [8] states that a property has a coarse threshold if and only if it is 'local' in the sense that it correlates with the property of containing a subset of a bounded size. Friedgut's criterion, and Bourgain's formulation [8, Appendix] that extends it to a more general setting, have been an instrumental tool in proving that various properties have a sharp threshold.

Even though more than twenty years have passed since the work of Friedgut was published, only a handful of Ramsey properties have been shown to (or not to) have a sharp threshold: First, Friedgut and Krivelevich [11] showed that for any tree T (bar stars) and for any number of colours r (except for $r = 2$ in the case where T is the path of length three), the property that any r-colouring of the edges of $G_{n,p}$ contains a monochromatic copy of T has a sharp threshold. Next, Friedgut, Rödl, Ruciński, and Tetali [12] established sharpness of the threshold for the corresponding property of the triangle, but only in the case where the number of colours r is equal to two. Much later, Friedgut, Hàn, Person, and Schacht [10] proved sharpness of the threshold in the context of van der Waerden's theorem, again only in the two-colour case. Building on ideas from [10], Schacht and Schulenburg [28] returned to the setting of

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Ramsey's theorem and managed to extend the result of [12] from triangles to all nearly bipartite graphs (see below) whereas Schulenburg [29] showed sharpness of the threshold in the context of Schur's theorem; both these results apply only to the two-colour case.

The main result of this paper is a common generalisation of all the above works, save for [11]. We view Ramsey properties of random subsets as statements about non-r-colourability of subhypergraphs that these random subsets induce in the hypergraph $\mathcal{H}$ that represents copies of a given structure A in the ground set V. Our main result supplies sufficient conditions on a sequence of uniform hypergraphs that guarantee that non-r-colourability of the random induced subhypergraph $\mathcal{H}[V_p]$, and thus the corresponding Ramsey property of random sets, has a sharp threshold. We postpone the exact statement of our theorem to Section 2 and, in the remainder of this section, present several interesting corollaries of this general result.

1.1. Graph properties

Given graphs G and H and an integer $r \geqslant 2$, we write $G \to (H)_r$ if every r-colouring of the edges of G contains a monochromatic copy of H. For the vast majority of pairs H and r, the location of the threshold for the property $G_{n,p} \to (H)_r$ is determined by a simple parameter of H, called the *2-density*, defined by

$$m_2(H) \coloneqq \max\left\{\frac{e_F - 1}{v_F - 2} : \emptyset \neq F \subseteq H\right\} \cup \left\{\frac{1}{2}\right\}.$$

The following statement was proved in a series of papers of Rödl and Ruciński [23, 24, 25]. (The necessity for the special treatment of paths of length three in the case $r = 2$, originally missed by Rödl and Ruciński, was noticed by Friedgut and Krivelevich [11].)

Theorem 1.1 [25]. *Let $r \geqslant 2$ be an integer and suppose that H is a nonempty graph whose at least one component is not a star or (in the case $r = 2$) a path of length three. There exist positive constants c_0 and c_1 such that*

$$\lim_{n\to\infty} \mathbb{P}\big(G_{n,p} \to (H)_r\big) = \begin{cases} 1 & \text{if } p \geqslant c_1 \cdot n^{-1/m_2(H)}, \\ 0 & \text{if } p \leqslant c_0 \cdot n^{-1/m_2(H)}. \end{cases}$$

In other words, Theorem 1.1 states that, for most pairs H and r, the function $n^{-1/m_2(H)}$ is a threshold for the property $G_{n,p} \to (H)_r$. In the case where H is a tree, Friedgut and Krivelevich [11] gave a complete characterisation of those pairs for which the corresponding threshold is coarse (when H is a star or when $r = 2$ and H is a path of length three) or sharp (all other pairs H and r). Deciding the sharpness of the threshold for the property $G_{n,p} \to (H)_r$ turned out to be much harder in the case where H contains a cycle. Here, our knowledge is only fragmentary. The monumental work of Friedgut, Rödl, Ruciński, and Tetali [12] established sharpness of the threshold in the case where H is the triangle and $r = 2$ using a very elaborate, long, and technical argument. The authors of [12] speculated that the threshold is sharp whenever H contains a cycle, for any number of colours, but so far this has been confirmed only when H is nearly bipartite[1] and strictly 2-balanced[2] and $r = 2$ in the recent work of Schacht and Schulenburg [28].

We prove that the threshold is sharp for a much broader family of graphs that includes all cliques and for any number of colours. We call a graph H *collapsible* if, for every edge e of H and every endpoint a of e, there is an edge f of H and a homomorphism from $H \setminus f$ to $H \setminus e$ that maps both endpoints of f to a.[3] It is not difficult to verify (see Section 7.1) that every graph that is either complete or nearly bipartite is collapsible. Unfortunately, not every graph is collapsible; for example, the Petersen graph is not collapsible (see Appendix C).

[1] A graph H is *nearly bipartite* if $\chi(H \setminus e) \leqslant 2$ for some edge e of H.
[2] A graph H is *strictly 2-balanced* if $m_2(F) < m_2(H)$ for every strict, nonempty subgraph $F \subseteq H$.
[3] Given a graph H and an edge e, we denote by $H \setminus e$ the graph with same vertex set as H whose edge set is $E(H) \setminus \{e\}$.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Theorem 1.2. *Suppose that H is a strictly 2-balanced, collapsible graph that is not a forest and $r \geqslant 2$ is an integer. There exist positive constants c_0 and c_1 and a function $c(n)$ satisfying $c_0 \leqslant c(n) \leqslant c_1$ such that, for every positive ε,*

$$\lim_{n\to\infty} \mathbb{P}\big(G_{n,p} \to (H)_r\big) = \begin{cases} 1 & \text{if } p \geqslant (1+\varepsilon)c(n) \cdot n^{-1/m_2(H)}, \\ 0 & \text{if } p \leqslant (1-\varepsilon)c(n) \cdot n^{-1/m_2(H)}. \end{cases}$$

Remark. In fact, when $r = 2$, we may replace the assumption that H is collapsible with a seemingly weaker assumption that H is semi-collapsible (see Definition 7.8). However, we did not find an example of a graph that is semi-collapsible and not collapsible.

It would be extremely interesting to extend Theorem 1.2 to a broader class of graphs as well as to verify whether or not the function c from the statement of the theorem has a limit as $n \to \infty$.

1.2. Arithmetic properties

We say that a set Y of elements of some ambient additive group is *r-Schur*, for some integer $r \geqslant 2$, and write that $Y \in \mathcal{S}_r$ if every r-colouring of the elements of Y admits a monochromatic sum, by which we mean three *distinct* elements $a, b, c \in Y$ such that $a + b = c$, all coloured the same way. Schur's theorem [30] states that, for any fixed r, the set $[\![N]\!] := \{1, \dots, N\}$ is r-Schur whenever N is sufficiently large. Similarly, given integers $k \geqslant 3$ and $r \geqslant 2$ and a set Y of elements of some additive group, we say that Y is *(k, r)-van der Waerden* and write $Y \in \mathcal{W}(k, r)$ if every r-colouring of the elements of Y admits a monochromatic k-term arithmetic progression. The well-known theorem of van der Waerden [31] states that, for all k and r, the set $[\![N]\!]$ is (k, r)-van der Waerden provided that N is sufficiently large (as a function of k and r).

Rödl and Ruciński [25] proved that, for any $k \geqslant 3$ and any number of colours $r \geqslant 2$, the function $N^{-1/(k-1)}$ is a threshold for the property $\mathcal{W}(k, r)$ in the set $[\![N]\!]_p$. Soon afterwards, Graham, Rödl, and Ruciński [14] showed that the function $N^{-1/2}$ is a threshold for the property $\mathcal{S}_r$, for any $r \geqslant 2$. Friedgut, Hàn, Person, and Schacht [10] showed that the former threshold is sharp whereas Schulenburg [29] proved the analogous statement for the latter threshold. Both of these results are valid only for random subsets of the cyclic group $\mathbb{Z}_N$ and, crucially, only in the case $r = 2$. Even though the results of [14, 25] are established for random subsets of $[\![N]\!]$, their proofs can be easily adapted to yield analogous statements for subsets of $\mathbb{Z}_N$. (In fact, the 1-statements in the nonmodular setting imply the 1-statements in the modular setting. As for the 0-statements, the results presented in Section 1.3 below generalise and strengthen both these results.)

We establish sharpness of the thresholds for $\mathcal{W}(k, r)$ and $\mathcal{S}_r$ in random subsets of $\mathbb{Z}_N$ (in the case of Schur's theorem, we additionally require N to be prime) for all $k \geqslant 3$ and all $r \geqslant 2$. As in [28, 29], the reason for replacing $[\![N]\!]$ with $\mathbb{Z}_N$ is that our approach requires the ground set to have a transitive group of symmetries that preserves the structure defining the property (Schur triples or k-APs).

Theorem 1.3. *For all integers $k \geqslant 3$ and $r \geqslant 2$, there are constants $c_0 \leqslant c_1$ and a function $c_0 \leqslant c(N) \leqslant c_1$ such that for all $\varepsilon > 0$,*

$$\lim_{N\to\infty} \mathbb{P}\big((\mathbb{Z}_N)_p \in \mathcal{W}(k, r)\big) = \begin{cases} 1, & p \geqslant (1+\varepsilon)c(N) \cdot N^{-1/(k-1)}, \\ 0, & p \leqslant (1-\varepsilon)c(N) \cdot N^{-1/(k-1)}. \end{cases}$$

Theorem 1.4. *For any integer $r \geqslant 2$, there are constants $c_0 \leqslant c_1$ and a function $c_0 \leqslant c(N) \leqslant c_1$ such that for all $\varepsilon > 0$,*

$$\lim_{\substack{N\to\infty \\ \text{is prime}}} \mathbb{P}\big((\mathbb{Z}_N)_p \in \mathcal{S}_r\big) = \begin{cases} 1, & p \geqslant (1+\varepsilon)c(N) \cdot N^{-1/2}, \\ 0, & p \leqslant (1-\varepsilon)c(N) \cdot N^{-1/2}. \end{cases}$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

1.3. List Ramsey problems

A main new theme in our analysis that paves the way to proving sharp threshold results in the case where the number of colours is larger than two is a list-colouring generalisation of the Ramsey problem. The main result of this work views Ramsey results, such as Ramsey's theorem, van der Waerden's theorem, or Schur's theorem mentioned above, as statements about non-r-colourability of certain hypergraphs. In the proof of this result, however, we encounter the more general problem of *list colouring* a hypergraph from a given assignment of lists (of size two) to its vertices, which can be viewed as a list Ramsey problem. Let us mention that a list-colouring variant of Ramsey's theorem that is closely related to the one considered here was recently introduced by Alon, Bucić, Kalvari, Kuperwasser, and Szabó [1] and subsequently studied by Fox, He, Luo, and Xu [6].

In this section, we consider threshold phenomena associated with such list Ramsey problems in the context of van der Waerden's and Schur's theorems. We say that a set Y of elements of an additive group is *list-Schur* if there exists an assignment of two-element lists to the elements of Y such that every colouring of the elements of Y with colours from their lists must admit a monochromatic sum. We define the notion of a *list-k-van der Waerden* sets analogously. Note that every 2-Schur (resp. $(2,k)$-van der Waerden) set is also list-Schur (resp. list-van der Waerden), but the converse is not necessarily true. Our arguments yield, with very little extra work, the following strengthenings of the 0-statements of the aforementioned results of [14, 25] that establish the location of the threshold for van der Waerden's and Schur's theorems in random sets of integers.

Theorem 1.5. *For every integer $k \geqslant 3$, there is a constant c such that, for every sequence X of sets of elements of an additive group such that $|X| \to \infty$ and every $p \leqslant c \cdot |X|^{-1/(k-1)}$,*

$$\mathbb{P}\big(X_p \text{ is list-}k\text{-van der Waerden}\big) \to 0.$$

Theorem 1.6. *There is a constant c such that, for every sequence X of sets of elements of an additive group such that $|X| \to \infty$ and every $p \leqslant c \cdot |X|^{-1/2}$,*

$$\mathbb{P}\big(X_p \text{ is list-Schur}\big) \to 0.$$

In fact, both Theorems 1.5 and 1.6 are straightforward consequences of the following more general statement on the 2-choosability (i.e., 2-list-colourability) threshold of almost-linear hypergraphs, whose short (two and a half pages) proof is given in Section 7.4.

Theorem 1.17. *Suppose that $s \geqslant 3$ and that a sequence of s-uniform hypergraphs $\mathcal{H}$ satisfies $\Delta_2(\mathcal{H}) = O(1)$ and $v(\mathcal{H}) \to \infty$. There is a positive c such that, for every $p \leqslant c \cdot v(\mathcal{H})^{-1/(s-1)}$,*

$$\mathbb{P}\big(\mathcal{H}[V_p] \text{ is 2-choosable}\big) \to 1.$$

1.4. Organisation

The rest of the paper is organised as follows. In Section 2, we introduce our general theorem, which gives sufficient conditions (which we discuss in detail) on a sequence of hypergraphs that guarantee a sharp threshold for the property of r-colourability. Section 3 offers an outline of the proof of this general theorem. At the end of that section, we formulate two key statements that imply our result.

The bulk of the work is spent proving these statements. Section 4 provides some external tools, such as the sharp threshold criterion and the hypergraph container lemma, which we then spend some time honing to our needs. Subsequently, we prove both statements in Sections 5 and 6.

Finally, having wrapped up the proof of the main result, Section 7 turns to applying the theorem in the various settings we mentioned previously: for graphs, arithmetic progressions, and Schur triples.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

2. The main result

As we have mentioned above, we will view Ramsey properties of random sets as statements about non-r-colourability of random hypergraphs. Given a hypergraph $\mathcal{H}$ with vertex set V and a real $p \in [0,1]$, we will denote by $\mathcal{H}_p$ the subhypergraph of $\mathcal{H}$ induced by the random set V_p. If the edges of $\mathcal{H}$ are all copies of a structure A in a set V (for example, the edge sets of all copies of a graph H in K_n), then non-r-colourability of $\mathcal{H}_p$ is equivalent to the random set V_p having the corresponding r-colour Ramsey property with respect to A (in our example, the property $G_{n,p} \to (H)_r$). Since we are interested in threshold phenomena, we will almost always consider infinite sequences of hypergraphs whose sizes tend to infinity.

Our main result supplies a sufficient condition on a sequence $\mathcal{H}$ of uniform hypergraphs that guarantees that the property that $\mathcal{H}_p$ is not r-colourable has a sharp threshold. This sufficient condition is a conjunction of five assumptions. We first give a brief overview of these five assumptions and state our result and return to discussing them in detail in the remainder of this section.

2.1. Overview

Suppose that $\mathcal{H}$ is a sequence of s-uniform hypergraphs and let $r \geqslant 2$ be an integer. The following function takes centre stage in our considerations:

$$p_{\mathcal{H}} := \left(\frac{v(\mathcal{H})}{e(\mathcal{H})}\right)^{1/(s-1)}.$$

In order to phrase the five assumptions on the sequence $\mathcal{H}$ that guarantee that non-r-colourability has a sharp threshold in $\mathcal{H}_p$, we need to introduce three simple notions. A *star* in $\mathcal{H}$ is a collection of $r-1$ edges that pairwise intersect in a single vertex called the centre of the star. A *constellation* is a collection of s disjoint stars whose centres form an edge of $\mathcal{H}$. A star formed by edges $A_1, \dots, A_{r-1}$ and centred at v is *rainbow* if there are distinct colours $i_1, \dots, i_{r-1} \in [\![r]\!]$ such that, for each j, all vertices of $A_j \setminus \{v\}$ are coloured i_j. A constellation is rainbow if its s constituent stars are rainbow and have the same colour pattern (the set $\{i_1, \dots, i_{r-1}\}$). The conjunction of the following five assumptions implies that non-r-colourability has a sharp threshold in $\mathcal{H}_p$:

(A1) *Symmetry.* The hypergraph $\mathcal{H}$ is *symmetric* in the sense that its group of automorphisms $\mathrm{Aut}(\mathcal{H})$ acts transitively on the vertex set of $\mathcal{H}$.

(A2) *Nonclusteredness.* The hypergraph $\mathcal{H}$ is *nonclustered*, which means (roughly speaking) that $\mathcal{H}$ satisfies the assumptions of the hypergraph container lemma with density parameter $p_{\mathcal{H}}$. (See Section 2.2.)

(A3) *Weak threshold.* The function $p_{\mathcal{H}}$ is a threshold for the property that $\mathcal{H}_p$ is not r-colourable.

(A4) *Choosability of typical bounded-sized subsets.* The random set $V(\mathcal{H})_{p_{\mathcal{H}}}$ a.a.s. does not contain any set W with $O(1)$ vertices for which $\mathcal{H}[W]$ is not choosable from 2-element lists of colours in $[\![r]\!]$. (See Section 2.3.)

(A5) *The rainbow star-constellation property.* Every partial r-colouring of the vertices of $\mathcal{H}$ that makes a constant proportion of its stars rainbow must make a constant proportion of its constellations rainbow as well. (See Section 2.4.)

Theorem 2.1. *Let $s \geqslant 3$ and $r \geqslant 2$ be integers and let $\mathcal{H}$ be a sequence of s-uniform hypergraphs. If $\mathcal{H}$ satisfies assumptions (A1)–(A5), then there exists a function $\hat{p} = \Theta(p_{\mathcal{H}})$ such that the following holds for every positive ε:*

$$\mathbb{P}\big(\mathcal{H}_p \text{ is } r\text{-colourable}\big) \to \begin{cases} 1 & \text{if } p \leqslant (1-\varepsilon)\hat{p}, \\ 0 & \text{if } p \geqslant (1+\varepsilon)\hat{p}. \end{cases}$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Verifying assumptions (A1) and (A2) for our applications of the theorem will be completely straightforward. Assumption (A3) is not at all easy to check, but, for the three applications of the main theorem we consider in this work, it had been established by earlier works. Moreover, it is now standard to derive the 1-statement in (A3) from assumption (A2) and a property we term *robust noncolourability*, see Section 2.5. Verifying assumption (A4), which is closely related to establishing the 0-statement in (A3), takes the most effort. Assumption (A5) holds trivially when every set of $\Omega(v(\mathcal{H}))$ vertices induces $\Omega(e(\mathcal{H}))$ edges, which is the case in the context of van der Waerden's theorem and Ramsey's theorem for bipartite graphs. In the two remaining applications of the theorem discussed here – Schur's theorem and Ramsey's theorem for nonbipartite, collapsible graphs – establishing this assumption requires a nontrivial argument. Finally, let us mention that there are natural sequences of hypergraphs, for which one would expect non-r-colourability to have a sharp threshold, that satisfy assumptions (A1)–(A4), but fail to satisfy (A5), see Appendix C.

2.2. Nonclusteredness

Given a hypergraph $\mathcal{H}$ and a set $T \subseteq V(\mathcal{H})$, we will denote by $\deg_{\mathcal{H}}(T)$ the degree of T in $\mathcal{H}$, that is,

$$\deg_{\mathcal{H}}(T) \coloneqq |\{A \in \mathcal{H} : T \subseteq A\}|.$$

Further, for an integer $t \geqslant 1$, we let $\Delta_t(\mathcal{H})$ be the maximum degree of a t-element set of vertices, defined by

$$\Delta_t(\mathcal{H}) \coloneqq \max\{\deg_{\mathcal{H}}(T) : T \subseteq V(\mathcal{H}) \text{ and } |T| = t\}.$$

We are now ready to define the notion of nonclusteredness from assumption (A2).

Definition 2.2. A sequence of nonempty, s-uniform hypergraphs $\mathcal{H}$ is called *nonclustered* if

$$\Delta_1(\mathcal{H}) = O\left(\frac{e(\mathcal{H})}{v(\mathcal{H})}\right) \qquad \text{and} \qquad \Delta_t(\mathcal{H}) \ll p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \quad \text{for } t \in \{2, \dots, s-1\}.$$

Fact 2.3. Suppose that $\mathcal{H}$ is a nonclustered sequence of s-uniform hypergraphs.

(i) We have

$$\Delta_s(\mathcal{H}) = 1 = p_{\mathcal{H}}^{s-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}.$$

(ii) If $s \geqslant 3$, then $1 \leqslant \Delta_{s-1}(\mathcal{H})/\Delta_s(\mathcal{H}) \ll p_{\mathcal{H}}^{-1}$, implying that $p_{\mathcal{H}} \to 0$.

Fact 2.4. The following sequences of hypergraphs are nonclustered:

- The hypergraph of (the edge sets of) copies of a strictly 2-balanced graph in K_n.
- The hypergraph of k-term arithmetic progressions in the cyclic group $\mathbb{Z}_N$.
- The hypergraph of Schur triples in any Abelian group.

2.3. Choosability of typical bounded-sized subsets

Recall that a hypergraph $\mathcal{G}$ is *2-choosable* from a set C of colours if, for every assignment $L\colon V(\mathcal{G}) \to \binom{C}{2}$ of 2-element lists of colours to the vertices of $\mathcal{G}$, there exists a proper colouring of $\mathcal{G}$ that assigns to each vertex $v \in V(\mathcal{G})$ a colour from its list L_v. Given a hypergraph $\mathcal{H}$ and integers $k \geqslant 1$ and $r \geqslant 2$, define

$$\mathcal{N}_k(\mathcal{H}) \coloneqq \big\{W \subseteq V(\mathcal{H}) : |W| \leqslant k \text{ and } \mathcal{H}[W] \text{ is not 2-choosable from } [\![r]\!]\big\}.$$

The precise statement of assumption (A4) is that, for every $k \geqslant 1$,

$$\mathbb{P}\big(V(\mathcal{H})_{p_{\mathcal{H}}} \supseteq W \text{ for some } W \in \mathcal{N}_k(\mathcal{H})\big) \to 0. \tag{1}$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

In fact, our argument may require that (1) holds also when we replace $p_{\mathcal{H}}$ with some $p = \Theta(p_{\mathcal{H}})$. Fortunately, these two statements are completely equivalent, see Lemma 4.6.

Finally, it is worth pointing out that the assumption on choosability of typical bounded-sized subsets is necessary for non-2-colourability of $\mathcal{H}_p$ to have a sharp threshold at some $\hat{p} = \Theta(p_{\mathcal{H}})$. Indeed, if (1) fails for some constant k, then the probability that $V(\mathcal{H})_p$ contains some $W \in \mathcal{N}_k(\mathcal{H})$, which clearly makes $\mathcal{H}_p$ not 2-colourable, is bounded away from zero for every $p = \Omega(p_{\mathcal{H}})$.

2.4. The rainbow star-constellation property

Our final assumption (A5) has a much less obvious connection to the problem at hand, but it conveniently fits into our framework.

Definition 2.5. A collection $A_1, \dots, A_k$ of edges of a hypergraph $\mathcal{H}$ is called a *k-star* (or simply a *star*) if there exists a vertex v of $\mathcal{H}$ such that $A_i \cap A_j = \{v\}$ for every pair of distinct $i, j \in [\![k]\!]$; the vertex v is called the *centre* of the star and $A_1 \cup \dots \cup A_k$ is called the *support* of the star.

Definition 2.6. A collection of stars with pairwise-disjoint supports whose centres form an edge of $\mathcal{H}$ is called a *constellation*. The edge induced by the centres of the stars forming a constellation is called the *base* of the constellation.

Definition 2.7. Suppose that some vertices of a hypergraph $\mathcal{H}$ are coloured with the elements of $[\![r]\!]$, for some integer $r \geqslant 2$, and let $i \in [\![r]\!]$ be an arbitrary colour. We say that an $(r-1)$-star $\{A_j\}_{j \in [\![r]\!] \setminus \{i\}}$ centred at v is *i-rainbow* if, for every $j \in [\![r]\!] \setminus \{i\}$, all vertices of $A_j \setminus \{v\}$ are coloured j. A constellation is *i-rainbow* if all stars comprising it are *i*-rainbow. Finally, a star/constellation is *rainbow* if it is *i*-rainbow for some $i \in [\![r]\!]$.

A fairly straightforward calculation (Lemma 6.3) shows that every nonclustered sequence of s-uniform hypergraphs $\mathcal{H}$ contains $\Theta\big(e(\mathcal{H})^{r-1}/v(\mathcal{H})^{r-2}\big)$ many $(r-1)$-stars and $\Theta\big(e(\mathcal{H})^{s(r-1)+1}/v(\mathcal{H})^{s(r-1)}\big)$ constellations of $(r-1)$-stars. We will say that such a sequence $\mathcal{H}$ has the rainbow star-constellation property for r colours if every partial r-colouring of the vertices of $\mathcal{H}$ that makes a constant proportion of all its $(r-1)$-stars rainbow also makes a constant proportion of all its constellations rainbow.

Definition 2.8. Given an integer $r \geqslant 2$ and a sequence of s-uniform hypergraphs $\mathcal{H}$, we say that $\mathcal{H}$ has the *rainbow star-constellation property* for r colours if every partial colouring of $V(\mathcal{H})$ with elements of $[\![r]\!]$ that induces $\Omega\big(e(\mathcal{H})^{r-1}/v(\mathcal{H})^{r-2}\big)$ rainbow stars must also induce $\Omega\big(e(\mathcal{H})^{s(r-1)+1}/v(\mathcal{H})^{s(r-1)}\big)$ rainbow constellations.

2.5. The weak threshold assumption

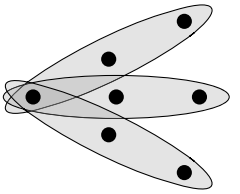
We conclude this section with a short discussion on how assumption (A3) might possibly be derived from (A2) and (A4) and yet another 'supersaturation' assumption on $\mathcal{H}$ that we term robust non-r-colourability.

Definition 2.9. Given an integer $r \geqslant 2$, we say that a sequence of hypergraphs $\mathcal{H}$ is *robustly non-r-colourable* if every r-colouring of the vertices of $\mathcal{H}$ makes a constant proportion of the edges of $\mathcal{H}$ monochromatic, that is, if every $c\colon V(\mathcal{H}) \to [\![r]\!]$ satisfies

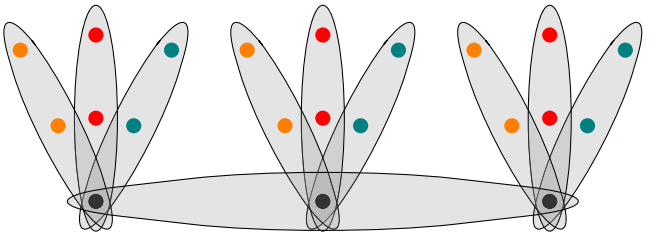
$$\sum_{i=1}^{r} e\big(\mathcal{H}[c^{-1}(i)]\big) = \Omega\big(e(\mathcal{H})\big).$$

The following facts can be derived from the corresponding Ramsey statements using simple averaging arguments and are thus considered folklore.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press



(A) A star



(B) A rainbow constellation

Figure 1. *Stars and constellations.*

Fact 2.10. The following sequences of hypergraphs are robustly non-r-colourable:

- The hypergraph of (the edge sets of) copies of a fixed nonempty graph in K_n.
- The hypergraph of k-term arithmetic progressions in the cyclic group $\mathbb{Z}_N$.
- The hypergraph of Schur triples in any Abelian group.

If a sequence of hypergraph is nonclustered and robustly non-r-colourable, then $O(p_{\mathcal{H}})$ is an upper bound on any threshold function of non-r-colourability. This fact can be shown by a straightforward adaptation of the argument of Nenadov and Steger [21], who showed that robust $(r+1)$-colourability of the sequence $\mathcal{H}$ of hypergraphs representing copies of a given graph H in K_n implies non-r-colourability of a typical $\mathcal{H}_p$ for all $p \gg n^{-1/m_2(H)}$. (See also [3, Section 8] for a slightly different version of this argument that shows the exact statement of the proposition below.) For the sake of completeness, we recreate this argument in Appendix A.

Proposition 2.11. *Let $r \geqslant 2$ and $s \geqslant 2$ be integers. For every nonclustered, robustly non-r-colourable sequence $\mathcal{H}$ of s-uniform hypergraphs, there exists a constant C such that, for every $p \geqslant Cp_{\mathcal{H}}$,*

$$\mathbb{P}\big(\mathcal{H}_p \text{ is } r\text{-colourable}\big) \leqslant \exp\big(-\Omega\big(p \cdot v(\mathcal{H})\big)\big).$$

Remark. Proposition A.2 proves a more general version of this statement. It is also shown in the appendix (in the remark following this proposition) that if $s \geqslant 3$, then $p \cdot v(\mathcal{H}) \to \infty$ when $p = \Omega(p_{\mathcal{H}})$.

It would be worth looking into the following problem, motivated by [21, Lemma 6] and [20, Meta-Theorem].

Problem 2.12. Does assumption (A4), with $r = 2$, imply that $\Omega(p_{\mathcal{H}})$ is a lower bound on the threshold for r-colourability of every nonclustered sequence of s-uniform hypergraphs, provided that $s \geqslant 3$?

One could consider strengthening the assumption of being nonclustered by further assuming that, for each $t \in \{2, \dots, s-1\}$, the inequality $\Delta_t(\mathcal{H}) \ll p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}$ hides some polynomial (in $v(\mathcal{H})$) factor. We remark that the three families of sequences of hypergraphs from the statement of Fact 2.4 all enjoy such strengthened nonclusteredness property. It is not unlikely that one can extend the methods from

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

subsection 7.4 and from the subsequent works [17, 18] to prove the statement under this strengthened assumption.

3. An outline of the proof

Assume that $r \geqslant 2$ and $s \geqslant 3$ and suppose that $\mathcal{H}$ is a sequence of s-uniform hypergraphs that satisfies assumptions (A1)–(A5). For the sake of brevity, we will denote the vertex set of $\mathcal{H}$ by V, its cardinality by N, and write that a set $W \subseteq V$ is r-colourable if and only if the induced subhypergraph $\mathcal{H}[W]$ is. Finally, assume to the contrary that the threshold for V_p not being r-colourable is coarse.

3.1. Boosters and dichotomy

Our point of departure will be Friedgut's criterion, in Bourgain's formulation, which tells us that there is a positive constant c, a sequence p satisfying

$$c \leqslant \mathbb{P}\big(V_p \text{ is not } r\text{-colourable}\big) \leqslant 1 - c, \tag{2}$$

and a family $\mathbb{B}$ of constant-sized subsets of V such that

$$\mathbb{P}\big(\exists B \in \mathbb{B} \text{ s.t. } B \subseteq V_p\big) > c \tag{3}$$

and every $B \in \mathbb{B}$ is a *booster*.

Definition 3.1. Given $\delta > 0$ and $p \in [0, 1]$, a set $B \subseteq V$ is called a (p, δ)-*booster* if

$$\mathbb{P}\big(V_p \text{ is not } r\text{-colourable} \mid B \subseteq V_p\big) > \mathbb{P}\big(V_p \text{ is not } r\text{-colourable}\big) + \delta.$$

Observe that assumption (A3) and (2) imply that $p = \Theta(p_{\mathcal{H}})$. Using the symmetry assumption (A1), we will expand on Friedgut's criterion and show that a typical sample $Z \sim V_p$ exemplifies a sort of dichotomy, in the following precise sense.

Step I. There are constants $\alpha, \varepsilon > 0$, an integer K, and $p = \Theta(p_{\mathcal{H}})$ such that, for any family $\mathcal{F} \subseteq \binom{\mathcal{H}}{\leqslant K}$ with

$$\mathbb{P}\big(\exists B \in \mathcal{F} \text{ s.t. } B \subseteq V_p\big) < \alpha, \tag{4}$$

there is a set $B_0 \in \binom{V(\mathcal{H})}{\leqslant K} \setminus \mathcal{F}$ such that the following holds. For infinitely many values N, the set $Z \sim V_p$ satisfies the following with probability larger than α:

$$\mathbb{P}(Z \cup h(B_0) \text{ is not } r\text{-colourable} \mid Z) > \alpha,$$

where h is taken u.a.r. from the set of symmetries of $\mathcal{H}$, and

$$\mathbb{P}\big(Z \cup V_{\varepsilon p} \text{ is not } r\text{-colourable} \mid Z\big) \leqslant 1/2.$$

Remark. Note that the second part of the dichotomy, stating that the probability that $Z \cup V_{\varepsilon p}$ is not r-colourable is strictly less than one, implies that Z must be r-colourable.

In other words, we are guaranteed the existence of two sets, B_0 and Z, with two properties that seem at odds. While a positive proportion of the symmetric copies of the constant-sized B_0 *interact* with Z – that is, Z ceases to be r-colourable once we add them – the probability that the random set $V_{\varepsilon p}$ interacts with Z is bounded away from one. We will call these interacting symmetric copies of B_0 *activated* boosters.

Furthermore, we are allowed to trim some undesirable properties from both B_0 and Z. In the case of B_0, this can be done by requiring that $B_0 \notin \mathcal{F}$ whereas in the case of Z, this can be done as $Z \sim V_p$ satisfies the assertion with probability bounded away from zero. We should remark at this point that

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

the family $\mathcal{F}$ we are going to choose will be *symmetric*, that is, if $B \in \mathcal{F}$, then $h(B) \in \mathcal{F}$ for every $h \in \mathrm{Aut}(\mathcal{H})$. Therefore, if $B_0 \notin \mathcal{F}$, then the same is true for every other symmetric copy of it.

Our aim is to use these two statements to get a contradiction. Specifically, we will show that, with a suitable choice of properties for B_0 and Z that exploit assumptions (A2)–(A5), the existence of many activated boosters implies that $Z \cup V_{\varepsilon p}$ is not r-colourable with probability arbitrarily close to one. The methodology of the argument will be very much in tune with the previous works of Friedgut, Hán, Person, and Schacht [10], Schacht and Schulenburg [28], and Schulenburg [29]. However, in order to extend these results we require some novel ideas. Adding the choosability assumption (A4) to the mix will allow us to argue for sharpness in three or more colours. At the same time, the introduction of the rainbow star-constellation property (A5) empowers the argument to target more complex combinatorial structures. As we will soon see, certain aspects of the argument simplify if the Turán problem that is associated with the structure is degenerate (as is the case with arithmetic progressions or bipartite graphs). The previous methods were able to argue about such cases, and even go further to tackle structures that are nearly degenerate (e.g., Schur triples and nearly bipartite graphs). Property (A5) encodes a weaker sense of degeneracy that generalises the previous cases and also applies for new structures, such as cliques of any size.

We will present our proof in two rounds. The first round will be a (spoiler alert) failed attempt, which will still show in essence how to utilise the assumption about choosability of typical subsets of bounded size (which we will enforce on B_0 via an appropriate choice of $\mathcal{F}$) to gain structural information on proper r-colourings of Z. The second round will address the breaking point of that approach, a very large union bound over r-colourings of Z, and remedy it using the Hypergraph Container Lemma of Saxton and Thomason [27] and also of Balogh, Morris, and Samotij [2]. This approach will, in turn, require us to strengthen one of the claims made in the first round.

3.2. First attempt

We start with Step (I) and get Z and a family of interacting boosters, all of which are symmetric copies of B_0. Define the hypergraph $\mathcal{B}$ on the vertex set V whose edges are all symmetric copies of B_0, that is, all $h(B_0)$ with $h \in \mathrm{Aut}(\mathcal{H})$. Since $\mathcal{H}$ is symmetric, $\mathcal{B}$ is regular and, therefore,

$$\Delta_1(\mathcal{B}) = |B_0| \cdot \frac{e(\mathcal{B})}{v(\mathcal{H})} \leqslant K \cdot N^{-1} \cdot e(\mathcal{B}).$$

We will impose some structural assumptions on B_0. It is natural to require that B_0 is r-colourable, since otherwise $Z \cup h(B_0)$ would be not r-colourable with probability one, which would in turn suggest that the threshold is actually coarse. (This was the only assumption on B_0 imposed in previous works [10, 28, 29].) We will go one step further. Instead of ensuring that B_0 is only r-colourable, we will make use of assumption (A4) and require that it is 2-choosable from lists in $[\![r]\!]$. We may do so as (A4) implies that the family $\mathcal{F}$ comprising all non-2-choosable subsets of V with at most K vertices satisfies the condition (4) in Step (I). Note that this property is symmetric, so requiring it from B_0 guarantees that it is fulfilled by all $B \in \mathcal{B}$.

Let $\mathcal{B}_Z \subseteq \mathcal{B}$ comprise only the copies of B_0 that interact with Z, that is,

$$\mathcal{B}_Z \coloneqq \big\{B \in \mathcal{B} : Z \cup B \text{ is not } r\text{-colourable}\big\}.$$

The first assertion of Step (I) translates to $e(\mathcal{B}_Z) > \alpha \cdot e(\mathcal{B})$. One of the desirable properties of Z would allow us to find a subfamily $\mathcal{B}'_Z \subseteq \mathcal{B}_Z$ of our activated boosters – satisfying $e(\mathcal{B}'_Z) \geqslant \alpha/2 \cdot e(\mathcal{B})$ – whose members interact with Z in a very well-behaved manner. First, if $B \in \mathcal{B}'_Z$, then B and Z are disjoint. Further, suppose that some activated booster $B \in \mathcal{B}_Z$ does not intersect Z. Since both Z and B are r-colourable, the fact that $Z \cup B$ is not means that there is an edge of $\mathcal{H}[Z \cup B]$ that intersects both B and Z. Call the set of all such edges the *interface* between B and Z. The second desirable property, which we may impose using assumption (A2), is that, if $B \in \mathcal{B}'_Z$, then each edge in the interface between B and Z has exactly one vertex in B (and the remaining vertices in Z).

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Definition 3.2. Given $U \subseteq V$, let $\mathrm{Col}(U)$ denote the set of all proper r-colourings of U, that is, colourings that do not admit a monochromatic edge.

Definition 3.3. A set $U \subseteq V$ *threatens* a vertex $v \in V$ if there are $u_1, \dots, u_{s-1} \in U$ such that $\{u_1, \dots, u_{s-1}, v\}$ is an edge of $\mathcal{H}$.

Let $\psi \in \mathrm{Col}(Z)$ be a proper colouring of Z. We say that ψ *forces* a vertex $v \in V$ to the colour $i \in [\![r]\!]$ if v is threatened by $\psi^{-1}(j)$ for every $j \neq i$. (Equivalently, ψ forces v to the colour i if and only if v is the centre of an i-rainbow star.) The motivation behind the definition is the following fact: If ψ forces v to the colour i, then every proper colouring of $Z \cup \{v\}$ that extends ψ must assign the colour i to v. Let $F_i(\psi)$ denote the set of vertices that ψ forced to the colour i and let $F(\psi) \coloneqq F_1(\psi) \cup \cdots \cup F_r(\psi)$ be the set of forced vertices.

Step II (first attempt). There is some $\lambda > 0$ such that $|F(\psi)| \geqslant \lambda N$ for all $\psi \in \mathrm{Col}(Z)$.

We prove this statement in two steps. First, we show that in order to force a constant fraction of the vertices, it is enough to force at least one vertex in a constant fraction of the activated boosters. Second, we show that every colouring $\psi \in \mathrm{Col}(Z)$ forces a vertex inside every activated booster.

Lemma 3.4 (Many forced boosters $\to$ many forced vertices). *If F is a set which intersects a constant fraction of the activated boosters, then $|F| = \Omega(N)$.*

Proof. Let $\beta > 0$ be some constant such that F intersects $\beta \cdot e(\mathcal{B}'_Z)$ activated boosters. Since each vertex belongs to at most $\Delta_1(\mathcal{B})$ boosters, we have

$$\beta \cdot \alpha/2 \cdot e(\mathcal{B}) \leqslant \beta \cdot e(\mathcal{B}'_Z) \leqslant |F| \cdot \Delta_1(\mathcal{B}) \leqslant |F| \cdot K \cdot N^{-1} \cdot e(\mathcal{B}),$$

which implies that F has $\Omega(N)$ elements. □

Lemma 3.5 (Activated booster $\to$ forced booster). *Every $\psi \in \mathrm{Col}(Z)$ forces at least one vertex in every booster.*

Proof. Suppose that this is not true and there is a booster $B \in \mathcal{B}'_Z$ whose every vertex is not threatened by at least two colour classes of ψ. Since B is 2-choosable, we would be able to find a colouring $\phi \in \mathrm{Col}(B)$ using only these nonthreatening colours. However, since $Z \cup B$ is not r-colourable, there must be an edge $\{u_1, \dots, u_{s-1}, b\} \in \mathcal{H}$, where $u_1, \dots u_{s-1} \in Z$ and $b \in B$, such that $\psi(u_1) = \cdots = \psi(u_{s-1}) = \phi(b)$. (Indeed, since $B \in \mathcal{B}'_Z$, every edge in the interface between B and Z has exactly one vertex in B.) This would mean that, contrary to our assumption, $\psi^{-1}(\phi(b))$ threatens b. □

By the pigeonhole principle, one of the forced sets, say $F_i(\psi)$, has $\Omega(N)$ vertices. We would like to show that this set induces $\Omega(e(\mathcal{H}))$ edges. (This is essentially equivalent to $\Omega(e(\mathcal{H}))$ edges being the base edge of an i-rainbow constellation.) Such a claim is always true when $\mathcal{H}$ describes a 'degenerate' structure; e.g., if $\mathcal{H}$ is the hypergraph of copies of a bipartite graph or when $\mathcal{H}$ is the hypergraph of arithmetic progressions. Unfortunately, some hypergraphs of interest (e.g., the hypergraph of copies of any non-bipartite graph) contain independent sets of cardinality $\Omega(N)$. However, using the assumption that $\mathcal{H}$ has the rainbow star-constellation property, we will be able to argue that a.a.s. $Z \sim V_p$ has the property that, for every $\psi \in \mathrm{Col}(Z)$, every large set $F_i(\psi)$ must induce many edges. With foresight, we state a stronger version of this property that extends also to *partial* r-colourings of Z.

Step III. For any $\beta > 0$, there exists a $\gamma > 0$ such that a.a.s. $Z \sim V_p$ has the following property: For any partial r-colouring ψ of Z and any $i \in [\![r]\!]$, if $|F_i(\psi)| \geqslant \beta N$, then $F_i(\psi)$ induces at least $\gamma \cdot e(\mathcal{H})$ edges.

The proof of this statement will appear in Section 6. We will just mention that it employs the Hypergraph Container Lemma to transfer the supersaturation statement given by the rainbow star-constellation assumption into the sparse regime.

Note now that if $e \subseteq F_i(\psi)$ is an edge, then ψ cannot be extended to e while staying proper. This is because the elements of $F_i(\psi)$ must all be coloured i, but this makes e monochromatic. The fact that

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

$F_i(\psi)$ contains $\Omega(e(\mathcal{H}))$ edges, together with the assumption that the hypergraph $\mathcal{H}$ is not clustered, makes it extremely unlikely that such an edge will fail to appear in $V_{\varepsilon p}$. The following estimate follows from Janson's inequality (Theorem 4.9).

Lemma 3.6. *Suppose that $A \subseteq V$ induces $\Omega(e(\mathcal{H}))$ edges. The probability that $V_{\varepsilon p}$ avoids all edges of A is bounded from above by* $\exp(-\Omega(pN))$.

Thus far, we have demonstrated that the probability that a *given* $\psi \in \mathrm{Col}(Z)$ could be extended to $Z \cup V_{\varepsilon p}$ is bounded from above by $\exp(-\Omega(pN))$. However, in order to get the desired contradiction to the first assertion of Step (I), we would like to show that the probability that *some* proper colouring of Z can be extended to $V_{\varepsilon p}$ tends to zero. To this end, let us take the union bound over all proper colourings. Alas, the only bound we have at our disposal is $|\mathrm{Col}(Z)| = \exp(O(|Z|)) = \exp(O(pN))$, which is not good enough. Moreover, we cannot significantly improve the upper bound established by Lemma 3.6, since the set $V_{\varepsilon p}$ is empty with probability approximately $\exp(-\varepsilon pN)$.

3.3. Second attempt

Since the breaking point of our first attempt was the union bound over all proper colourings of Z, we will try to make this union bound more efficient by excluding many colourings in $\mathrm{Col}(Z)$ at once. To this end, note that if ψ_0 were a partial colouring of Z that forced $\Omega(N)$ vertices to some colour, we could still apply Step (III) and Lemma 3.6 to learn that the probability that ψ_0 can be extended to $V_{\varepsilon p}$ is at most $\exp(-\Omega(pN))$. Moreover, if ψ_0 cannot be extended to $V_{\varepsilon p}$, then neither can any proper colouring $\psi \in \mathrm{Col}(Z)$ that extends ψ_0. It thus suffices to find a family of $\exp(o(|Z|)) = \exp(o(pN))$ partial colourings of Z, each forcing a constant proportion of the vertices to some colour, such that every element of $\mathrm{Col}(Z)$ is an extension of some member of the family.

As was hinted before, we will employ the Container Lemma [2, 27] to find such a family. To this end, we will identify every colouring $\psi\colon Z \to [\![r]\!]$ with the set $\{(z, \psi(z))\colon z \in Z\}$ and define a hypergraph $\mathcal{T}$ on the vertex set $Z \times [\![r]\!]$ such that every $\psi \in \mathrm{Col}(Z)$ will correspond to an independent set of $\mathcal{T}$. The Container Lemma will allow us to construct a family $\mathcal{C}$ of subsets of $Z \times [\![r]\!]$, which we will call *containers*, such that:

(C1) Any independent set of $\mathcal{T}$, and thus every $\psi \in \mathrm{Col}(Z)$, is contained in some $C \in \mathcal{C}$.
(C2) Every container $C \in \mathcal{C}$ induces only a small fraction of the edges of $\mathcal{T}$.

We will say that a set $C \subseteq Z \times [\![r]\!]$ is a *restricted colouring* if every $z \in Z$ has at least one 'available' colour in C, that is, if $(z, i) \in C$ for some $i \in [\![r]\!]$. Note that a set $C \subseteq Z \times [\![r]\!]$ that is not a restricted colouring cannot contain any colouring of Z. Given a restricted colouring $C \subseteq Z \times [\![r]\!]$, we define its *determined colouring* ψ_C to be the maximal partial colouring agreed upon by all the colourings contained by C. (That is, $\psi_C(z) = i$ if and only if i is the unique colour such that $(z, i) \in C$.) Note that (C1) implies that every $\psi \in \mathrm{Col}(Z)$ extends the determined colouring ψ_C of some $C \in \mathcal{C}$. The crux of our argument is showing that one may define $\mathcal{T}$ in such a way that condition (C2) implies that, for every $C \in \mathcal{C}$, the determined colouring ψ_C forces many vertices to some colour, so that we can apply Step (III).

The precise definition of the hypergraph $\mathcal{T}$ is somewhat technical, but the idea behind it is fairly straightforward. Given a booster $B \in \mathcal{B}'_Z$, we will say that colourings $\psi\colon Z \to [\![r]\!]$ and $\phi\colon B \to [\![r]\!]$ are *consistent* if no edge in the interface between B and Z is monochromatic under $\psi \cup \phi$. Since each such interface edge has exactly one vertex in B (by the definition of $\mathcal{B}'_Z$), two colourings ψ and ϕ are consistent if and only if no vertex $b \in B$ is threatened by $\psi^{-1}(\phi(b))$. A key observation is that the fact that two colourings ψ and ϕ as above are consistent always has a small 'witness' set in $Z \times [\![r]\!]$. Indeed, for a given ϕ, one can certify that ψ is consistent with ϕ by specifying the values of ψ on vertices of Z that belong to edges of the interface between B and Z. Such minimal 'witness' sets for all $B \in \mathcal{B}'_Z$ and all proper colourings $\phi \in \mathrm{Col}(B)$ are the edges of our hypergraph $\mathcal{T}$. The fact that every $B \in \mathcal{B}'_Z$ is an activated booster means that no two colourings $\psi \in \mathrm{Col}(Z)$ and $\phi \in \mathrm{Col}(B)$ are consistent and, consequently, every $\psi \in \mathrm{Col}(Z)$ is an independent set of $\mathcal{T}$.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

We say that a restricted colouring C *inter-activates* a booster $B \in \mathcal{B}'_Z$ if every proper colouring $\phi \in \mathrm{Col}(B)$ is inconsistent with each colouring of Z contained in C. Our definition of $\mathcal{T}$ guarantees that a restricted colouring C will induce an edge in $\mathcal{T}$ for every booster $B \in \mathcal{B}'_Z$ which it fails to inter-activate. (This edge will be a witness to some pair $\phi \in \mathrm{Col}(B)$ and $\psi \subseteq C$ being consistent.) Since every container $C \in \mathcal{C}$ induces only a small proportion of all edges of $\mathcal{T}$, it must therefore inter-activate all but a small fraction of all boosters in $\mathcal{B}'_Z$. Finally, an argument similar to the one used in the proof of Lemma 3.5 shows that if C inter-activates a booster B, then ψ_C must force at least one vertex in B to some colour. (A key insight here is realising that the property that C inter-activates a booster depends only on the determined colouring ψ_C.) This implies, by Lemma 3.4, that ψ_C must force $\Omega(N)$ vertices.

Our final concern is that the family $\mathcal{C}$ of containers will have at most $\exp(o(|Z|))$ elements. In order to guarantee this, we will need to demonstrate some control over the edge set of $\mathcal{T}$. In fact, it will be sufficient to bound the largest size of an edge of $\mathcal{T}$ by an absolute constant and to show that $\Delta_1(\mathcal{T}) \cdot |Z| = O(e(\mathcal{T}))$ and $\Delta_2(\mathcal{T}) \cdot |Z| \ll e(\mathcal{T})$. Luckily, this will be possible yet again by trimming further undesirable properties from Z in Step (I). For example, the expected number of interacting edges between V_p and B is bounded by some constant and consequently the witness sets will also be constant sized. Using this and other related properties of Z and the booster family, we will be able to show that the hypergraph $\mathcal{T}$ is 'well-behaved', which will allow us to use the Container Lemma to derive the following.

Step II (refined). There is a family $\mathcal{C}$ of $\exp(o(|Z|))$ subsets of $Z \times [\![r]\!]$ with the following properties:

(1) Every proper colouring of Z is an extension of the determined colouring ψ_C of some container $C \in \mathcal{C}$.
(2) For every container $C \in \mathcal{C}$, the determined colouring ψ_C forces $\Omega(N)$ vertices.

Our union bound argument is now brought back to life and the proof is finally settled.

3.4. Finalising the argument

Our goal for the next three sections is to tie up the loose ends of the proof, by filling in the gaps in the outline presented above. Recall that we want to show that, assuming non-r-colourability does not have a sharp threshold in $\mathcal{H}_p$, the following statements hold for $Z \sim V_p$ with probability bounded away from zero:

(I) For some constant $\varepsilon > 0$, the hypergraph $\mathcal{H}\big[Z \cup V_{\varepsilon p}\big]$ is r-colourable with probability at least $1/2$.
(II) Each proper r-colouring of Z extends one of $\exp(o(|Z|))$ partial colourings, each of which forces $\Omega(N)$ elements in V to some colour.
(III) Each of our partial colourings admits a set of $\Omega(e(\mathcal{H}))$ edges whose all vertices are forced to the same colour.

As we argued above, this would lead to a contradiction. Indeed, each individual partial colouring from the family cannot be extended to any set that contains one of the $\Omega(e(\mathcal{H}))$ 'forced' edges. By Lemma 3.6, the probability that $V_{\varepsilon p}$ omits all these edges is at most $\exp\big(-\Omega(pN)\big)$. A union bound over all partial colourings yields that $\mathcal{H}\big[Z \cup V_{\varepsilon p}\big]$ is not r-colourable with probability tending to one.

We prove statements (I), (II), and (III) over the course of Sections 4, 5, and 6. More precisely, Section 4 contains the formal statements of Friedgut's sharp threshold criterion, along with a routine corollary thereof (Step (I) in the proof outline), as well as the statements of the Hypergraph Container Lemma and Janson's inequality, of which Lemma 3.6 is a simple corollary.

Next, in Section 5, we show how the abundance of boosters can be used to construct the efficient family of partial colourings (the refined version of Step (II) in the proof outline). In order to formally phrase this result using the notions from assumption (A5), we shift the terminology from forced vertices to (centres of) rainbow stars. The result is the following theorem.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Theorem 3.7. *Suppose that $s \geqslant 3$ and $r \geqslant 2$ and let $\mathcal{H}$ be a sequence of s-uniform hypergraphs that satisfies assumptions (A1)–(A4). If non-r-colourability does not have a sharp threshold in $\mathcal{H}_p$, then there exist a constant $\varepsilon > 0$ and a subsequence $p = \Theta(p_{\mathcal{H}})$ such that, letting $Z \sim V(\mathcal{H})_p$, the following holds with probability at least ε:*

(i) *The hypergraph $\mathcal{H}\big[Z \cup V(\mathcal{H})_{\varepsilon p}\big]$ is r-colourable with probability at least $1/2$.*
(ii) *There exists a family Ψ of $\exp\big(o(|Z|)\big)$ partial $[\![r]\!]$-colourings of Z such that*
  (a) *every proper $[\![r]\!]$-colouring of $\mathcal{H}[Z]$ extends some partial colouring in Ψ and*
  (b) *for every colouring in Ψ, some $\Omega\big(v(\mathcal{H})\big)$ vertices of $\mathcal{H}$ are centres of rainbow stars.*

Finally, Section 6 provides a proof of the following sparse analogue of the rainbow star-constellation property (Step (III) in the proof outline). We remark again that this theorem is trivial if the hypergraph $\mathcal{H}$ is 'degenerate' in the sense that every subset of $\Omega(N)$ vertices induces $\Omega(e(\mathcal{H}))$ edges. As a result, we are spared from a bulk of the proof when $\mathcal{H}$ is the hypergraph of copies of a bipartite graph or the hypergraph of arithmetic progressions of a prescribed length.

Theorem 3.8. *Suppose that $s \geqslant 3$, $r \geqslant 2$, and $\varepsilon > 0$. Let $\mathcal{H}$ be a sequence of s-uniform hypergraphs that satisfies assumptions (A2) and (A5) and suppose that $Z \sim V(\mathcal{H})_p$ for some $p = \Theta(p_{\mathcal{H}})$. With probability at least $1 - \varepsilon$, for every partial colouring ψ of Z with elements of $[\![r]\!]$, if $\Omega\big(v(\mathcal{H})\big)$ vertices of $\mathcal{H}$ are centres of rainbow stars, then $\Omega\big(e(\mathcal{H})\big)$ edges of $\mathcal{H}$ are bases of rainbow constellations.*

Theorems 3.7 and 3.8, supplemented with Lemma 3.6, imply Theorem 2.1.

4. Preliminaries and tools

4.1. Coarse thresholds

The first piece of machinery that we require is a characterisation of properties with coarse thresholds. In the proof outline, this was captured by Step (I), which we aim to formalise and prove here. In general, Friedgut's work [8] and Bourgain's subsequent extension of it [8, Appendix] provide a criterion for the appearance of a sharp threshold. The criterion states that every property that fails to have a sharp threshold must correlate with a 'local' property, that is, the appearance of a bounded-sized subset. In other words, if a property $\mathcal{P}$ has a coarse threshold, then there is another, local, property $\mathcal{P}'$ – with the same threshold as $\mathcal{P}$ – such that both are positively correlated.

For our purpose, we use a reformulation of Bourgain's aforementioned result, which appears in [9]. We will introduce a relevant definition and then state the theorem.

Definition 4.1 (Boosters). Suppose that $\mathcal{P}$ is a property of subsets of a finite set V. Given a $p \in [0,1]$ and a positive number δ, we call a set $B \subseteq V$ a (p,δ)-*booster* if

$$\mathbb{P}(V_p \cup B \in \mathcal{P}) \geqslant \mathbb{P}(V_p \in \mathcal{P}) + \delta.$$

Theorem 4.2 (The Sharp Threshold Criterion). *For all positive α and C, there exist positive δ, η, p_0, and K such that the following holds. Suppose that $\mathcal{P}$ is a monotone property of subsets of a finite set V and, for each $p \in [0,1]$, let $\mu(p) \coloneqq \mathbb{P}(V_p \in \mathcal{P})$. If, for some $0 < p \leqslant p_0$*

$$\alpha \leqslant \mu(p) \leqslant 1 - \alpha \qquad \text{and} \qquad \mu'(p) \leqslant C/p,$$

then there is a family $\mathcal{B} \subseteq \binom{V}{\leqslant K}$ satisfying

$$\mathbb{P}(B \subseteq V_p \text{ for some } B \in \mathcal{B}) > \eta$$

such that every $B \in \mathcal{B}$ is a (p,δ)-booster for $\mathcal{P}$.

We will now use the criterion to derive a general result in the spirit of Step (I) of the proof outline. First, it guarantees the existence of boosters with 'typical' characteristics. Second, it asserts the following

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

dichotomy for properties $\mathcal{P}$ with coarse thresholds: while there are many (constant-sized) boosters whose addition to V_p lands us immediately in $\mathcal{P}$, adding to V_p a random set of density εp does not increase the probability of being in $\mathcal{P}$ substantially. We again introduce a definition and move on to stating and proving the result.

Definition 4.3. We say that a set $B \subseteq V$ is an *active booster* for a set $Z \subseteq V$ if $Z \cup B \in \mathcal{P}$.

Proposition 4.4. *Let $\mathcal{P}$ be a nontrivial, monotone property of subsets of a (sequence of) finite set(s) V. For each $p \in [0,1]$, let $\mu(p) \coloneqq \mathbb{P}(V_p \in \mathcal{P})$ and let $\hat{p} \coloneqq \mu^{-1}(1/2)$. If $\mathcal{P}$ does not have a sharp threshold and $\hat{p} = o(1)$, then there exist positive constants δ, ε, and K and an infinite subsequence $p = \Theta(\hat{p})$ such that:*

(1) *For every family $\mathcal{F}$ of subsets of V satisfying*

$$\mathbb{P}\big(V_p \supseteq W \text{ for some } W \in \mathcal{F}\big) < \varepsilon,$$

there is a (p,δ)-booster in $\binom{V}{\leqslant K} \setminus \mathcal{F}$.

(2) *For any family $\mathcal{B}$ of (p,δ)-boosters, letting $Z \sim V_p$, the following holds with probability at least ε:*
   (a) $\mathbb{P}(Z \cup V_{\varepsilon p} \in \mathcal{P} \mid Z) \leqslant 1/2$ *and*
   (b) *at least $\varepsilon|\mathcal{B}|$ elements of $\mathcal{B}$ are active boosters for Z.*

Proof. Suppose that $\mathcal{P}$ does not have a sharp threshold. This means that there are constants $c_1 < c_2$ and $\alpha > 0$ such that

$$\alpha \leqslant \mu(c_1 \cdot \hat{p}) \leqslant \mu(c_2 \cdot \hat{p}) \leqslant 1 - \alpha$$

on some infinite subsequence. Let $C \coloneqq 4c_2/(c_2 - c_1)$.

Claim 4.5. *For every $\gamma > 0$, there is a $p \in (c_1 \cdot \hat{p}, c_2 \cdot \hat{p})$ such that*

(a) $\mu'(p) \leqslant C/p$ *and*
(b) $\mu(p + \gamma p) - \mu(p) \leqslant C \cdot \gamma$

Proof. Fix an arbitrary γ' satisfying $0 < \gamma' < (c_2 - c_1)/2$. Since $0 \leqslant \mu \leqslant 1$, we have

$$\int_{c_1}^{c_2-\gamma'} \big(\mu((x+\gamma') \cdot \hat{p}) - \mu(x \cdot \hat{p})\big)\, dx \leqslant \int_{c_2-\gamma'}^{c_2} \mu(x \cdot \hat{p})\, dx \leqslant \gamma'.$$

In particular, there must be a p' satisfying $c_1 \cdot \hat{p} < p' < (c_2 - \gamma') \cdot \hat{p}$ such that

$$\mu(p' + \gamma' \cdot \hat{p}) - \mu(p') \leqslant \frac{\gamma'}{c_2 - c_1 - \gamma'} \leqslant \frac{2\gamma'}{c_2 - c_1}.$$

Further, as μ is increasing, there must be a $p \in (p', p' + (\gamma'/2) \cdot \hat{p}) \subseteq (c_1 \cdot \hat{p}, c_2 \cdot \hat{p})$ with

$$\mu'(p) \leqslant \frac{\mu(p' + (\gamma'/2) \cdot \hat{p}) - \mu(p')}{(\gamma'/2) \cdot \hat{p}} \leqslant \frac{4}{(c_2 - c_1) \cdot \hat{p}} \leqslant \frac{4c_2}{(c_2 - c_1) \cdot p} = \frac{C}{p},$$

as required. Finally, we show that the value p chosen above also satisfies inequality (b). If $\gamma \geqslant 1/C$, then the inequality is vacuous. Otherwise, if $\gamma < 1/C$, we let $\gamma' = 2c_2\gamma$. Since $\gamma' < (c_2 - c_1)/2$, by the definition of C, we have

$$p + \gamma p \leqslant p + \gamma \cdot c_2 \hat{p} = p + (\gamma'/2) \cdot \hat{p} \leqslant p' + \gamma' \cdot \hat{p}.$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Since $p > p'$ and μ is increasing, we have

$$\mu(p+\gamma p) - \mu(p) \leqslant \mu(p' + \gamma' \cdot \hat{p}) - \mu(p') \leqslant \frac{2\gamma'}{c_2 - c_1} = C\gamma,$$

as desired. □

Let $\mathcal{B}$ be an arbitrary nonempty subset of the family of (p,δ)-boosters supplied by Theorem 4.2. Fix a small positive constant γ, let B be a uniformly chosen random element of $\mathcal{B}$, and, for each $Z \subseteq V$, define

$$f(Z) = \mathbb{P}(Z \cup B \in \mathcal{P}) \qquad \text{and} \qquad g(Z) = \mathbb{P}(Z \cup V_{\gamma p} \in \mathcal{P}).$$

Our goal is to show that, for some positive constant ε,

$$\mathbb{P}\big(f(V_p) \geqslant \varepsilon \text{ and } g(V_p) \leqslant 1/2\big) \geqslant \varepsilon.$$

To this end, note first that Markov's inequality and the definition of a (p,δ)-booster imply

$$\mathbb{P}\big(f(V_p) < \varepsilon\big) = \mathbb{P}\big(1 - f(V_p) > 1 - \varepsilon\big) \leqslant \frac{1 - \mathbb{E}[f(V_p)]}{1-\varepsilon} \leqslant \frac{1 - \mu(p) - \delta}{1-\varepsilon}.$$

Since $g(Z) = 1$ when $Z \in \mathcal{P}$ and since $V_p \cup V_{\gamma p}$ is stochastically dominated by $V_{p+\gamma p}$, we have, again by Markov's inequality,

$$\begin{aligned}\mathbb{P}\big(g(V_p) > 1/2\big) &= \mathbb{P}(V_p \in \mathcal{P}) + \mathbb{P}\big(g(V_p) - \mathbb{1}_{V_p \in \mathcal{P}} > 1/2\big)\\ &\leqslant \mu(p) + 2\big(\mathbb{E}[g(V_p)] - \mathbb{E}[\mathbb{1}_{V_p \in \mathcal{P}}]\big)\\ &\leqslant 2\mu(p+\gamma p) - \mu(p) \leqslant \mu(p) + 2C\gamma.\end{aligned}$$

Finally,

$$\begin{aligned}\mathbb{P}\big(f(V_p) \geqslant \varepsilon \text{ and } g(V_p) \leqslant 1/2\big) &\geqslant 1 - \mathbb{P}\big(f(V_p) < \varepsilon\big) - \mathbb{P}\big(g(V_p) > 1/2\big)\\ &\geqslant 1 - \frac{1-\mu(p)-\delta}{1-\varepsilon} - \mu(p) - 2C\gamma\\ &\geqslant \delta - \frac{\varepsilon}{1-\varepsilon} - 2C\gamma \geqslant \delta/2 \geqslant \varepsilon,\end{aligned}$$

where the last two inequalities hold provided that ε and γ are sufficiently small (as functions of δ and C only). □

Finally, we will make use of the following simple lemma, which formalises the fact that all 'local' properties (i.e., properties of containing a bounded-sized subset from a given family) have coarse thresholds.

Lemma 4.6 (Local coarseness). *Let $\mathcal{F}$ be a family of subsets of a set V, all of which have at most K elements, and let $\mu(p) := \mathbb{P}(\exists B \in \mathcal{F}\colon B \subseteq V_p)$. Then, for any $c \in (0,1)$, we have $\mu(cp) \geqslant c^K \mu(p)$.*

Proof. We couple V_p and V_{cp} by viewing V_{cp} as the random subset $(V_p)_c \subseteq V_p$. Therefore,

$$\frac{\mu(cp)}{\mu(p)} = \mathbb{P}(\exists B \in \mathcal{F}\colon B \subseteq (V_p)_c \mid \exists B \in \mathcal{F}\colon B \subseteq V_p) \geqslant c^K.$$ □

4.2. Containers

Our second main tool is a hypergraph container lemma for almost independent sets. The standard versions of the container lemma, proved independently by Balogh, Morris, and Samotij [2] and by Saxton and Thomason [27], assert that every uniform hypergraph $\mathcal{G}$ admits a relatively small collection

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

of *containers* for independent sets, that is, a family $\mathcal{C}$ of subsets of $V(\mathcal{G})$, each containing only few edges of $\mathcal{G}$, such that every independent set of $\mathcal{G}$ is contained in some member of $\mathcal{C}$. Here, we will require a strengthening of this result that supplies a small collection of containers for the larger family of almost independent sets. Such stronger version of the container lemma was proved by Saxton and Thomason. Since the precise phrasing of this result that best fits our framework, Theorem B below, differs somewhat from [27, Corollary 3.6], we include a short derivation of the former from the latter in Appendix B.

Theorem 4.7. *For every positive integer k and all positive reals ε and K, there exist an integer t and a positive real δ such that the following holds. Suppose that a nonempty k-uniform (multi)hypergraph $\mathcal{G}$ with vertex set V and a positive real τ satisfy*

$$\Delta_\ell(\mathcal{G}) \leqslant K\tau^{\ell-1} \cdot \frac{e(\mathcal{G})}{v(\mathcal{G})}$$

for every $\ell \in [\![k]\!]$. Then, there exists a function $f\colon \mathcal{P}(V)^t \to \mathcal{P}(V)$ with the following properties:

(i) *For every set $I \subseteq V$ satisfying $e(\mathcal{G}[I]) \leqslant \delta\tau^k e(\mathcal{G})$, there are $S_1, \dots, S_t \subseteq I$ with at most $\tau v(\mathcal{G})$ elements each such that $I \subseteq f(S_1, \dots, S_t)$.*
(ii) *For every $S_1, \dots, S_t \subseteq V$, the set $f(S_1, \dots, S_t)$ induces fewer than $\varepsilon e(\mathcal{G})$ edges in $\mathcal{G}$.*

4.3. Janson's inequality

The final auxiliary result required by our argument is the well-known concentration inequality of Janson, which gives strong upper bounds on the lower tail probabilities of random variables counting how many sets from a given collection are contained in a binomial random subset. The version of Janson's inequality stated below differs from [16, Theorem 1] in the choice of notation, but it is otherwise equivalent.

Definition 4.8. Suppose that $\mathbf{A} = (A_1, \dots, A_k)$ is a sequence of (not necessarily distinct) events. The *pseudo-variance* of $\mathbf{A}$ is

$$\mathrm{Var}'(\mathbf{A}) \coloneqq \sum_{i \sim j} \mathbb{P}(A_i \cap A_j),$$

where the sum ranges over all ordered pairs $i, j \in [\![k]\!]$ such that A_i and A_j are not independent.

Remark. Note that $\mathrm{Var}'(A_1, \dots, A_k)$ is an upper bound on the variance of $\mathbb{1}_{A_1} + \cdots + \mathbb{1}_{A_k}$, hence the name pseudo-variance.

Theorem 4.9 (Janson's inequality [16]). *Suppose that Ω is a finite set, let $B_1, \dots, B_k$ be a sequence of (not necessarily distinct) subsets of Ω, and let $R \sim \Omega_p$ for some $p \in (0,1)$. For each $i \in [k]$, let X_i be the indicator of the event A_i that $B_i \subseteq R$ and let $X \coloneqq \sum_i X_i$. Then, for any $0 \leqslant t \leqslant \mathbb{E}[X]$,*

$$\mathbb{P}\big(X \leqslant \mathbb{E}[X] - t\big) \leqslant \exp\left(-\frac{t^2}{2\mathrm{Var}'(A_1, \dots, A_k)}\right).$$

We finish this section by deriving the following generalisation of Lemma 3.6.

Lemma 4.10. *Suppose that $s \geqslant 2$ and let $\mathcal{H}$ be a sequence of s-uniform hypergraphs that satisfies assumption (A2). If $p = \Theta(p_{\mathcal{H}})$, then, for every $\mathcal{H}' \subseteq \mathcal{H}$ with $\Omega(e(\mathcal{H}))$ edges,*

$$\mathbb{P}\big(e(\mathcal{H}'_p) = 0\big) = \exp\big(-\Omega\big(p \cdot v(\mathcal{H})\big)\big).$$

Proof. Write V for $V(\mathcal{H})$ and X for the number of edges of $\mathcal{H}'_p = \mathcal{H}'[V_p]$. We bound the expectation of X from below as follows:

$$\mu \coloneqq \mathbb{E}[X] = p^s e(\mathcal{H}') = \Omega\big(p^s e(\mathcal{H})\big).$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Further, we bound the pseudo-variance of the sequence $\mathbf{A}$ of events $B \subseteq V_p$ for all edges $B \in \mathcal{H}'$ from above:

$$\mathrm{Var}'(\mathbf{A}) = \sum_{i=1}^{s} \sum_{\substack{B,B' \in \mathcal{H}' \\ |B \cap B'| = i}} \mathbb{P}(B \cup B' \subseteq V_p) = O\left(\sum_{i=1}^{s} p^{2s-i} \cdot e(\mathcal{H}) \cdot \Delta_i(\mathcal{H}) \right).$$

Since $p = \Theta(p_{\mathcal{H}})$ and $\mathcal{H}$ is nonclustered,

$$\frac{\mu^2}{\mathrm{Var}'(\mathbf{A})} = \Omega\left(\min_{i \in [\![s]\!]} \frac{p_{\mathcal{H}}^i \cdot e(\mathcal{H})}{\Delta_i(\mathcal{H})} \right) = \Omega\big(p_{\mathcal{H}} \cdot v(\mathcal{H})\big).$$

Thus, Janson's inequality gives us that $\mathbb{P}(X = 0) = \exp\big(-\Omega\big(p \cdot v(\mathcal{H})\big)\big)$. □

5. Containers for colourings

In this section, we prove the following theorem, which encapsulates Steps (I) and (II) from the proof outline.

Theorem3.7. *Suppose that $s \geqslant 3$ and $r \geqslant 2$ and let $\mathcal{H}$ be a sequence of s-uniform hypergraphs that satisfies assumptions (A1)–(A4). If non-r-colourability does not have a sharp threshold in $\mathcal{H}_p$, then there exist a constant $\varepsilon > 0$ and a subsequence $p = \Theta(p_{\mathcal{H}})$ such that, letting $Z \sim V(\mathcal{H})_p$, the following holds with probability at least ε:*

(i) *The hypergraph $\mathcal{H}\big[Z \cup V(\mathcal{H})_{\varepsilon p}\big]$ is r-colourable with probability at least $1/2$.*
(ii) *There exists a family Ψ of $\exp\big(o(|Z|)\big)$ partial $[\![r]\!]$-colourings of Z such that*
  (a) *every proper $[\![r]\!]$-colouring of $\mathcal{H}[Z]$ extends some partial colouring in Ψ and*
  (b) *for every colouring in Ψ, some $\Omega\big(v(\mathcal{H})\big)$ vertices of $\mathcal{H}$ are centres of rainbow stars.*

5.1. Setup

Suppose that s, r, and $\mathcal{H}$ are as in the statement of the theorem. For the sake of brevity, throughout this section, we shall write V in place of $V(\mathcal{H})$. Assume that non-r-colourability does not have a sharp threshold in $\mathcal{H}_p$. Since $p_{\mathcal{H}}$ is a threshold function for this property, by assumption (A3), Proposition 4.4 supplies constants δ, ε, and K and an infinite sequence $p = \Theta(p_{\mathcal{H}})$ satisfying (1) and (2) in the proposition. Since $\mathcal{H}$ satisfies assumption (A4), invoking (1) in Proposition 4.4 with $\mathcal{F}$ being the family $\mathcal{N}_K(\mathcal{H})$ defined in Section 2.3, we obtain a (p, δ)-booster $B_0 \subseteq V$ of cardinality at most K such that $\mathcal{H}[B_0]$ is 2-choosable from $[\![r]\!]$.

Let $\mathcal{B}$ be the (multi)hypergraph on V whose edges are the images of B_0 via all automorphisms of $\mathcal{H}$. Since non-r-colourability is preserved under automorphisms of $\mathcal{H}$, every edge of $\mathcal{B}$ is also a (p, δ)-booster. Moreover, since $\mathcal{H}$ is symmetric, the hypergraph $\mathcal{B}$ is degree-regular and, consequently,

$$\Delta_1(\mathcal{B}) = |B_0| \cdot \frac{e(\mathcal{B})}{v(\mathcal{H})} \leqslant K \cdot \frac{e(\mathcal{B})}{v(\mathcal{H})}.$$

Let $\mathcal{Z}$ be the family of all subsets $Z \subseteq V$ such that $\mathcal{H}[Z \cup V_{\varepsilon p}]$ is r-colourable with probability at least $1/2$, but, for at least ε-proportion of $B \in \mathcal{B}$, the hypergraph $\mathcal{H}[Z \cup B]$ is not r-colourable. Property (2) in Proposition 4.4 states that $\mathbb{P}(V_p \in \mathcal{Z}) \geqslant \varepsilon$.

Finally, we define several constants. Let $K_{\mathcal{H}}$ be a constant satisfying

$$\Delta_1(\mathcal{H}) \leqslant K_{\mathcal{H}} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \qquad \text{and} \qquad \frac{p}{p_{\mathcal{H}}} \leqslant K_{\mathcal{H}}.$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Further, pick $\lambda = \varepsilon^2/8$ and let L be an integer satisfying

$$\frac{\left(K \cdot K_{\mathcal{H}}^s\right)^L}{L!} \leqslant \lambda. \tag{5}$$

5.2. *A recap of the proof outline*

Our goal now is to utilize the boosters in $\mathcal{B}$ in order to construct, for a typical $Z \in \mathcal{Z}$, a family Ψ of $\exp\left(o(|Z|)\right)$ partial colourings of Z that satisfies (a) and (b) from the statement of the theorem.

Given a $Z \in \mathcal{Z}$, let

$$\mathcal{B}_Z \coloneqq \left\{B \in \mathcal{B} : \mathcal{H}[Z \cup B] \text{ is not } r\text{-colourable}\right\}$$

be the family of boosters from $\mathcal{B}$ that are active for Z. By definition, no proper $[\![r]\!]$-colouring of Z can be extended to $Z \cup B$, for any $B \in \mathcal{B}_Z$. The assumption that $\mathcal{H}[B]$ is 2-choosable from $[\![r]\!]$ implies that, if a proper partial $[\![r]\!]$-colouring ψ_0 of Z cannot be extended to a booster in $B \in \mathcal{B}_Z$, then at least one vertex of B must be the centre of a rainbow star (equivalently, it must be forced to some colour), see Lemma 3.5. Therefore, it will be sufficient to make sure that our partial colourings do not extend to a constant proportion of the boosters in $\mathcal{B}_Z$.

We say that $A \in \mathcal{H}$ is an *interacting* edge between two sets if A is contained in their union and intersects both of them. Note that a proper partial colouring ψ_0 cannot be extended to B if and only if, for every proper colouring φ of B, some edge is monochromatic under $\psi_0 \cup \varphi$. This monochromatic edge cannot be fully contained in either Z or B, because both ψ_0 and φ are proper, and is therefore an interacting edge.

Define the *interface* of Z with B to be the set[4]

$$I(B, Z) \coloneqq \{A \setminus B : A \in \mathcal{H} \text{ is an interacting edge between } Z \text{ and } B\}.$$

We will say that a partial colouring ψ_0 of Z and a colouring φ of B are *consistent* if no edge that is interacting between Z and B is monochromatic under $\psi_0 \cup \varphi$. In this language, a partial colouring ψ_0 of Z can be extended to B precisely when ψ_0 is consistent with some proper colouring φ of B. Note that the interface $I(B, Z)$ contains all the information about ψ_0 that is needed to determine whether this is the case.

We say that $C \subseteq Z \times [\![r]\!]$ is a *restricted colouring* of Z if, for every $z \in Z$, there is some colour $i \in [\![r]\!]$ for which $(z, i) \in C$. Recall that, in this context, we identify colourings ψ of Z with the restricted colouring $\{(z, \psi(z)) : z \in Z\}$. Given a restricted colouring C, we define its *determined colouring* to be the partial colouring ψ_C, where $\psi_C(z) = i$ if every proper colouring $\psi \subseteq C$ has $\psi(z) = i$ (equivalently, if $i \in [\![r]\!]$ is the unique colour such that $(z, i) \in C$).

As we wrote in the outline, we will first define a hypergraph $\mathcal{T}$ on the vertex set $Z \times [\![r]\!]$ in such a way that proper colourings of Z will be independent in $\mathcal{T}$. We will then let the family Ψ of partial colourings be the determined colourings of a family of (restricted colourings that are) containers for the independent sets of $\mathcal{T}$. The edges of $\mathcal{T}$ will correspond to colourings of interfaces $I(B, Z)$, with $B \in \mathcal{B}_Z$, that are consistent with some proper colouring of B. In particular, we need to exercise control over the interfaces that Z has with the boosters in order to invoke the container lemma in an efficient way.

The rest of the section is organised as follows. We first prove that Z is disjoint from almost all boosters in $\mathcal{B}_Z$ and that the interface between Z and these boosters is bounded in size and well-behaved. Next, we will formally define the hypergraph $\mathcal{T}$ whose containers provide the family of partial colourings satisfying the first two properties. Finally, we will bound the degree sequence of $\mathcal{T}$ in order show that the family of containers, and therefore partial colourings, is small.

[4]This is a minor departure from the terminology used in the proof outline, where the interface between Z and B was the family of all edges A that are interacting between Z and B.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

5.3. Active boosters and interactions

Let $\mathcal{B}'_Z$ be the hypergraph obtained from $\mathcal{B}_Z$ by retaining only those $B \in \mathcal{B}_Z$ that satisfy all of the following:

(B1) The sets B and Z are disjoint.
(B2) Every nonempty set in $I(B,Z)$ has $s-1$ elements.
(B3) The sets in $I(B,Z)$ are pairwise disjoint.[5]
(B4) The family $I(B,Z)$ contains at most L sets.

The following two lemmas show that $\mathcal{B}_Z \setminus \mathcal{B}'_Z$ is small for a typical $Z \sim V_p$.

Lemma 5.1. *If $Z \sim V_p$, then the expected number of $B \in \mathcal{B}$ that fail one of (B1)–(B3) is $o\big(e(\mathcal{B})\big)$.*

Lemma 5.2. *If $Z \sim V_p$, then the expected number of $B \in \mathcal{B}$ that satisfy (B2) and (B3) but fail (B4) is at most $\lambda e(\mathcal{B})$.*

Finally, let

$$\mathcal{Z}' = \big\{Z \in \mathcal{Z} : e(\mathcal{B}'_Z) \geqslant (\varepsilon/2)\cdot e(\mathcal{B})\big\}.$$

By Lemmas 5.1 and 5.2 and since $\lambda = \varepsilon^2/8$, we have, for $Z \sim V_p$,

$$\begin{aligned}\mathbb{P}(Z \in \mathcal{Z}') &\geqslant \mathbb{P}(Z \in \mathcal{Z}) - \mathbb{P}\big(e(\mathcal{B}_Z \setminus \mathcal{B}'_Z) \geqslant (\varepsilon/2)\cdot e(\mathcal{B})\big)\\ &\geqslant \varepsilon - \frac{\mathbb{E}\big[e(\mathcal{B}_Z \setminus \mathcal{B}'_Z)\big]}{(\varepsilon/2)\cdot e(\mathcal{B})} \geqslant \varepsilon - 4\lambda/\varepsilon \geqslant \varepsilon/2.\end{aligned}$$

Proof of Lemma 5.1. Suppose that $Z \sim V_p$ and let X_1, X_2, and X_3 denote the numbers of sets B that fail conditions (B1), (B2), and (B3), respectively. We have

$$\mathbb{E}[X_1] = \sum_{B\in\mathcal{B}} \mathbb{P}(B \cap Z \neq \emptyset) \leqslant \sum_{B\in\mathcal{B}} \mathbb{E}[|B \cap Z|] = \sum_{v\in V} p \deg_{\mathcal{B}}(v) \leqslant v(\mathcal{H})\cdot p\Delta_1(\mathcal{B}) \ll e(\mathcal{B}),$$

since $p \ll 1$ and $\Delta_1(\mathcal{B}) = O\big(e(\mathcal{B})/v(\mathcal{H})\big)$. Since X_2 is at most the number of pairs $(A,B) \in \mathcal{H}\times\mathcal{B}$ such that $A \setminus B \subseteq Z$ and $1 \leqslant |A\setminus B| \leqslant s-2$, we have

$$\mathbb{E}[X_2] \leqslant \sum_{B\in\mathcal{B}} \sum_{t=1}^{s-2} \binom{|B|}{s-t} \cdot \Delta_{s-t}(\mathcal{H}) \cdot p^t \leqslant e(\mathcal{B}) \cdot \sum_{r=2}^{s-1} 2^K \cdot \Delta_r(\mathcal{H}) \cdot p_{\mathcal{H}}^{s-r},$$

where we used the fact that every edge of $\mathcal{B}$ contains at most K vertices. Since $\mathcal{H}$ is nonclustered, we have, for every $r \in \{2,\dots,s-1\}$,

$$\Delta_r(\mathcal{H}) \cdot p_{\mathcal{H}}^{s-r} \ll p_{\mathcal{H}}^{s-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} = 1,$$

implying that $\mathbb{E}[X_2] \ll e(\mathcal{B})$. Finally, we can bound $X_3 - X_2$ by the number of triples $(A, A', B) \in \mathcal{H}^2\times\mathcal{B}$ such that $A \cup A' \subseteq B \cup Z$, $1 \leqslant |(A\cap A')\setminus B| \leqslant s-2$, and $|A\cap B| = |A'\cap B| = 1$. We therefore get the following upper bound by counting the options for picking B, then A, then the size of the intersection $(A\cap A')\setminus B$, and finally A':

$$\mathbb{E}[X_3 - X_2] \leqslant \sum_{B\in\mathcal{B}} |B| \cdot \Delta_1(\mathcal{H}) \cdot \sum_{r=1}^{s-2} \binom{s-1}{r} \cdot |B| \cdot \Delta_{r+1}(\mathcal{H}) \cdot p_{\mathcal{H}}^{2s-2-r}.$$

[5] However, there could still be pairs of different $A, A' \in \mathcal{H}$ such that $A \setminus B = A' \setminus B \in I(B,Z)$.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Using our assumptions that every edge of $\mathcal{B}$ contains at most K vertices, $\Delta_1(\mathcal{H}) \leqslant O\big(e(\mathcal{H})/v(\mathcal{H})\big)$, and $\Delta_{r+1}(\mathcal{H}) \ll p_{\mathcal{H}}^r \cdot e(\mathcal{H})/v(\mathcal{H})$ for all $r \in \{1, \dots, s-2\}$, we have

$$\mathbb{E}[X_3 - X_2] \ll e(\mathcal{B}) \cdot p_{\mathcal{H}}^{2s-2} \cdot \left(\frac{e(\mathcal{H})}{v(\mathcal{H})}\right)^2 = e(\mathcal{B}).$$

The proof of the lemma is now complete. □

Proof of Lemma 5.2. For $B \in \mathcal{B}$ and $Z \subseteq V$, let $i(B, Z)$ be the largest integer ℓ such that there are $A_1, \dots, A_\ell \in \mathcal{H}$ satisfying:

(i) $A_1 \cup \cdots \cup A_\ell \subseteq B \cup Z$,
(ii) $|A_i \cap B| = 1$ for every $i \in [\![\ell]\!]$,
(iii) $A_1 \setminus B, \dots, A_\ell \setminus B$ are pairwise disjoint.

Observe that, if B satisfies (B2) and (B3), then $|I(B, Z)| = i(B, Z)$. In particular, the assertion of the lemma will follow if we show that $\mathbb{P}\big(i(B, Z) > L\big) \leqslant \lambda$.

To this end, for a positive integer ℓ, let $\mathcal{A}_\ell$ denote the collection of all sequences $A_1, \dots, A_\ell \in \mathcal{H}$ that satisfy (ii) and (iii) above and note that

$$|\mathcal{A}_\ell| \leqslant \big(|B| \cdot \Delta_1(\mathcal{H})\big)^\ell \leqslant \big(K \cdot \Delta_1(\mathcal{H})\big)^\ell.$$

Since, for every $(A_1, \dots, A_\ell) \in \mathcal{A}_\ell$, the set $(A_1 \cup \cdots \cup A_\ell) \setminus B$ has precisely $(s-1)\ell$ elements, we have

$$\mathbb{P}\big(A_1 \cup \cdots \cup A_\ell \subseteq B \cup Z\big) = p^{(s-1)\ell}.$$

Let X_ℓ be the number of sequences in $\mathcal{A}_\ell$ that satisfy (i). Since each such sequence has distinct coordinates and (i)–(iii) are invariant under any permutation of the sequence $(A_1, \dots, A_\ell)$, we may conclude that

$$\mathbb{P}\big(i(B, V_p) \geqslant \ell\big) \leqslant \frac{\mathbb{E}[X_\ell]}{\ell!} = \frac{|\mathcal{A}_\ell| \cdot p^{(s-1)\ell}}{\ell!} \leqslant \frac{\big(K \cdot \Delta_1(\mathcal{H}) \cdot p^{s-1}\big)^\ell}{\ell!}.$$

The assertion of the lemma follows because our assumptions imply that

$$K \cdot \Delta_1(\mathcal{H}) \cdot p^{s-1} \leqslant K \cdot K_{\mathcal{H}}^s \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \cdot p_{\mathcal{H}}^{s-1} = K \cdot K_{\mathcal{H}}^s,$$

see (5). □

5.4. *The hypergraph of transversals*

We will now formally define the hypergraph $\mathcal{T}$. Using the terminology of Section 5.2, our goal is the following: For every restricted colouring C, if its determined colouring ψ_C is consistent with some proper colouring of a booster $B \in \mathcal{B}'_Z$, then this fact will be evidenced by an edge of $\mathcal{T}$ inside C.

Let us discuss what that should mean. Suppose that φ is a proper colouring of a booster $B \in \mathcal{B}'_Z$. We will show that, for every restricted colouring C, its determined colouring ψ_C is consistent with φ if and only if there is a colouring $\psi_{C,\varphi} \colon \bigcup I(B, Z) \to [\![r]\!]$ that is consistent with φ and further satisfies $\psi_{C,\varphi} \subseteq C$. Since $\bigcup I(B, Z)$ contains at most $(s-1) \cdot L$ vertices, see (B2), these colourings $\psi_{C,\varphi}$ can serve as the kind of witnesses we are looking for. As a result, one could actually stop here, and take the edges of the hypergraph to be all such sets $\big\{(z, \psi_{C,\varphi}(z)) \colon z \in \bigcup I(B, Z)\big\}$. However, we will carry on and find witnesses that are minimal/irredundant.

To this end, suppose that an edge A interacting between B and Z is not monochromatic. Obviously, this means that two of its elements are coloured differently. However, we will use the fact that A intersects B in exactly one vertex (see (B2)) to conclude that either:

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

(1) ψ coloured two of the vertices of $A \setminus B$ in different colours, or
(2) ψ coloured some vertex of $A \setminus B$ in a colour different from $\varphi(A \cap B)$.

One might be tempted to ignore the first case, which seems superfluous. Indeed, even if two vertices are coloured in different colours, one of these colours must be different from the one in $\varphi(A \cap B)$. However, for a given $T \in I(B, Z)$, there can be many edges $A \in \mathcal{H}$ with $A \setminus B = T$ and picking a colour that is different from the one in $\varphi(A \cap B)$ for *every* such A might not always be possible.

We will define the edges of $\mathcal{T}$ using transversals of $I(B, Z)$. Given a proper colouring φ of B, we will pick one of two obstructions for every $T \in I(B, Z)$: either two coloured vertices (z, c), (z', c'), where $z, z' \in T$ and $c \neq c'$, or just one coloured vertex (z, c), where $z \in T$ and c different from the colour in $\varphi(A \cap B)$ for *every* $A \in \mathcal{H}$ such that $A \setminus B = T$.

We now define the hypergraph $\mathcal{T}$ formally. For every $(s-1)$-element set $T \subseteq Z$, let

$$N_{\mathcal{H}}(T; B) := \{b \in B : T \cup \{b\} \in \mathcal{H}\}$$

and note that $I(B, Z)$ comprises precisely those sets $T \subseteq Z \setminus B$ for which $|N_{\mathcal{H}}(T; B)| \geqslant 1$. In particular, given ψ and φ, there is an element $T \in I(B, Z)$ that ψ colours monochromatically from the colours in $\varphi\big(N_{\mathcal{H}}(T; B)\big)$.

Definition 5.3. We let $\mathcal{T}$ be the (multi)hypergraph with vertex set $Z \times [\![r]\!]$ whose (multi)set of edges is defined as follows. For every:

- active booster $B \in \mathcal{B}'_Z$,
- proper colouring $\varphi\colon B \to [\![r]\!]$ of $\mathcal{H}[B]$,
- set $I' \subseteq I(B, Z)$ satisfying $I' \supseteq \big\{T \in I(B, Z) : \varphi\big(N_{\mathcal{H}}(T; B)\big) = [\![r]\!]\big\}$, and
- disjoint transversals[6] $\{v_T : T \in I(B, Z)\}$ of $I(B, Z)$ and $\{v'_T : T \in I'\}$ of I', where $v_T \in T$ for every $T \in I(B, T)$ and $v'_T \in T$ for every $T \in I'$,

we add to $\mathcal{T}$ all edges of the form

$$\big\{(v_T, c_T) : T \in I(B, Z)\big\} \cup \big\{(v'_T, c'_T) : T \in I'\big\},$$

where:

- for each $T \in I(B, Z) \setminus I'$, the colour c_T is an arbitrary element of $[\![r]\!] \setminus \varphi\big(N_{\mathcal{H}}(T; B)\big)$;[7]
- for each $T \in I'$, the colours c_T and c'_T are two arbitrary, distinct elements of $[\![r]\!]$.

An immediate consequence of the definition should be that every proper colouring ψ of Z, when viewed as a restricted colouring, is an independent set of $\mathcal{T}$. This is because the determined colouring of ψ, which is just ψ itself, cannot be extended to any $B \in \mathcal{B}'_Z$.

Lemma 5.4. *Every proper colouring $\psi\colon Z \to [\![r]\!]$ of $\mathcal{H}[Z]$, viewed as a subset of $Z \times [\![r]\!]$, is an independent set in $\mathcal{T}$.*

Proof. Suppose that some proper colouring ψ of $\mathcal{H}[Z]$ contains an edge of $\mathcal{T}$. This means that there are an active booster $B \in \mathcal{B}'_Z$, a proper colouring φ of $\mathcal{H}[B]$, a set $I' \subseteq I(B, Z)$, and disjoint transversals $\{v_T : T \in I(B, Z)\}$ and $\{v'_T : T \in I'\}$ such that, for every $T \in I(B, Z)$, either $\psi(v_T) \notin \varphi\big(N_{\mathcal{H}}(T; B)\big)$ or $\psi(v_T) \neq \psi(v'_T)$. In particular, there is no $T \in I(B, Z)$ whose all elements receive the same colour from the set $\varphi\big(N_{\mathcal{H}}(T; B)\big)$. On the other hand, since B is disjoint from Z, see (B1), one may naturally define the colouring $\psi \cup \varphi\colon B \cup Z \to [\![r]\!]$. Since B is an active booster, this colouring $\psi \cup \varphi$ is not a proper colouring of $\mathcal{H}[Z \cup B]$, which means that there is an $A \in \mathcal{H}[Z \cup B]$ that $\psi \cup \varphi$ makes monochromatic. Since ψ and φ are proper colourings of $\mathcal{H}[Z]$ and of $\mathcal{H}[B]$, respectively, every such A must have a

[6] A *transversal* of a family $\mathcal{F}$ of sets is a subset X of $\bigcup \mathcal{F}$ such that $|X \cap F| = 1$ for every $F \in \mathcal{F}$.
[7] The choice of I' guarantees that the set $[\![r]\!] \setminus \varphi\big(N_{\mathcal{H}}(T; B)\big)$ is nonempty for every $T \in I(B, Z) \setminus I'$.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

nonempty intersection with both Z and B. Our choice of $\mathcal{B}'_Z$, see (B2), implies that $|A \cap B| = 1$ and, consequently, that $A \setminus B \in I(B, Z)$. Moreover,

$$(\psi \cup \varphi)(A) = \psi(A \setminus B) = \varphi(A \cap B),$$

which means that all elements of $A \setminus B$ receive the same colour as the unique vertex in $A \cap B$, which belongs to $N_{\mathcal{H}}(A \setminus B; B)$, a contradiction. □

More importantly, the definition implies that, for every restricted colouring C that induces a small number of edges in $\mathcal{T}$, its determined colouring ψ_C forces a constant fraction of the vertices of $\mathcal{H}$ to some colour. Recall that we denote the set of vertices that ψ_C forces to some colour by $F(\psi_C)$.

Lemma 5.5. *If $C \subseteq Z \times [\![r]\!]$ is a restricted colouring, then*

$$e(\mathcal{T}[C]) \geqslant e(\mathcal{B}'_Z) - |F(\psi_C)| \cdot \Delta_1(\mathcal{B}).$$

Proof. It suffices to find, for every active booster $B \in \mathcal{B}'_Z$ that does not contain any vertex of $F(\psi_C)$, an edge of $\mathcal{T}[C]$ that corresponds to B, as in Definition 5.3. For every $b \in B$, let $L(b) \subseteq [\![r]\!]$ be a list of some two colours that do not threaten b; there are at least two such colours since b is not forced to any colour. Let $\varphi\colon B \to [\![r]\!]$ be a proper colouring of $\mathcal{H}[B]$ from these lists; such a colouring exists by our assumption that $\mathcal{H}[B]$ is 2-choosable from $[\![r]\!]$, see (A4).

Fix an arbitrary $T \in I(B, Z)$ and let $N(T) = N_{\mathcal{H}}(T; B)$. If ψ_C colours all of T with a colour c_T, then c_T threatens every $b \in N(T)$ and hence $c_T \notin \bigcup_{b \in N(T)} L(b) \supseteq \varphi(N(T))$. In this case, we do not include T in I', we pick an arbitrary $v_T \in T$ and add $(v_T, c_T) \in C$ to the edge we construct. In the complementary case, either:

(i) ψ_C is not defined on all of T, which means that there is some vertex v in T and colours $c, c' \in [\![r]\!]$ such that $(v, c), (v, c') \in C$ (this is because $\pi_1(C) = Z$ and hence the only reason why $\psi_C(v)$ is not defined is that $|\{v\} \times [\![r]\!] \cap C| > 1$); or
(ii) ψ_C is defined on all of T but there are distinct vertices $v, v' \in T$ such that $\psi_C(v) \neq \psi_C(v')$.

Either way, since $|T| = s - 1 \geqslant 2$, there must be distinct $v_T, v'_T \in T$ and $c_T, c'_T \in [\![r]\!]$ such that $(v_T, c_T), (v'_T, c'_T) \in C$; we include T in I' and add (v_T, c_T) and (v'_T, c'_T) to the edge we construct. □

5.5. *The distribution of edges of $\mathcal{T}$*

The key step in the proof of Theorem 5 will be to construct a small family of containers for proper $[\![r]\!]$-colourings of $\mathcal{H}[Z]$ using the information provided by the hypergraph $\mathcal{T}$. In order to do this, we will first show that, for a typical choice of $Z \sim V_p$, we may find a subhypergraph $\mathcal{T}' \subseteq \mathcal{T}$ that contains almost as many edges as $\mathcal{T}$ and satisfies

$$\Delta_1(\mathcal{T}') = O\left(\frac{e(\mathcal{T}')}{|Z|}\right) \qquad \text{and} \qquad \Delta_2(\mathcal{T}') = o\left(\frac{e(\mathcal{T}')}{|Z|}\right).$$

We will then split $\mathcal{T}'$ into uniform hypergraphs $\mathcal{T}_1, \dots, \mathcal{T}_{2L}$, where, for each $u \in [\![2L]\!]$, the u-uniform hypergraph $\mathcal{T}_u$ comprises all edges of $\mathcal{T}'$ of cardinality u. Finally, we will apply the hypergraph container lemma, Theorem B, to each of the $\mathcal{T}_u$ in order to build containers for independent sets of $\mathcal{T}$: the container for an independent set I will be the intersection of the $2L$ containers for I in the hypergraphs $\mathcal{T}_1, \dots, \mathcal{T}_{2L}$. The assumptions on $\Delta_1(\mathcal{T}')$ and $\Delta_2(\mathcal{T}')$ will guarantee that there are only $\exp(o(|Z|))$ different containers. Lemma 5.5 provides a useful, structural description of all such containers.

In order to guarantee the existence of a $\mathcal{T}' \subseteq \mathcal{T}$ with the required properties, we shall estimate the ℓ^2-norm of the sequences of vertex degrees and vertex-pair degrees of $\mathcal{T}$ and invoke the following simple proposition.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Proposition 5.6. *Suppose that $\mathcal{G}$ is a (multi)hypergraph with vertex set V. For all positive integers t and m, there is a subhypergraph $\mathcal{G}' \subseteq \mathcal{G}$ with $e(\mathcal{G}') \geqslant e(\mathcal{G}) - m$ and*

$$\Delta_t(\mathcal{G}') \leqslant \frac{1}{m} \cdot \sum_{T \in \binom{V}{t}} \deg_{\mathcal{G}}(T)^2.$$

Proof. We may assume that $m \leqslant e(\mathcal{G})$; indeed, if $m > e(\mathcal{G})$, we may simply take $\mathcal{G}'$ to be the empty hypergraph. We obtain $\mathcal{G}'$ from $\mathcal{G}$ by iteratively removing m edges that contain some set $T \in \binom{V}{t}$ with largest degree. Since each of the m removed edges contained a t-element subset of degree at least $\Delta_t(\mathcal{G}')$, we have

$$\sum_{T \in \binom{V}{t}} \deg_{\mathcal{G}}(T)^2 \geqslant m \cdot \Delta_t(\mathcal{G}'),$$

as claimed. □

The following two crucial lemmas give upper bounds on the expectations of the ℓ^2-norms of the sequences of vertex degrees and vertex-pair degrees of $\mathcal{T}$.

Lemma 5.7. *If $Z \sim V_p$, then*

$$\mathbb{E}\left[\sum_{v \in V(\mathcal{T})} \deg_{\mathcal{T}}(v)^2\right] \leqslant \frac{\Gamma \cdot e(\mathcal{B})^2}{p_{\mathcal{H}} \cdot v(\mathcal{H})},$$

where Γ is a constant that depends only on K, $K_{\mathcal{H}}$, L, r, and s.

Lemma 5.8. *If $Z \sim V_p$, then*

$$\mathbb{E}\left[\sum_{T \in \binom{V(\mathcal{T})}{2}} \deg_{\mathcal{T}}(T)^2\right] \leqslant \frac{\sigma \cdot e(\mathcal{B})^2}{p_{\mathcal{H}} \cdot v(\mathcal{H})}$$

for some $\sigma = o(1)$.

Proof of Lemma 5.7. For every $v \in V$, define

$$\gamma(v) \coloneqq \left|\left\{A \in \mathcal{H} : v \in A \text{ and } |(A \setminus \{v\}) \cap Z| = s - 2\right\}\right|.$$

We claim that, for all $v \in Z$ and $i \in [\![r]\!]$,

$$\deg_{\mathcal{T}}(v, i) \leqslant \gamma(v) \cdot \Delta_1(\mathcal{B}) \cdot (rs)^{2L+K}, \tag{6}$$

where $\deg_{\mathcal{T}}$ counts edges with multiplicities. Indeed, if $v \in Z$, then $\gamma(v) \cdot \Delta_1(\mathcal{B})$ is an upper bound on the number of $B \in \mathcal{B}'_Z$ such that $v \in T$ for some $T \in I(B, Z)$. Every such B gives rise to at most

$$r^K \cdot \big((s-1)r\big)^L \cdot \sum_{i=0}^{L} \big((s-2)r\big)^i \leqslant (rs)^{2L+K}$$

edges of $\mathcal{T}$, counting multiplicities; indeed, there are at most r^K proper $[\![r]\!]$-colourings of B, at most $(s-1)r$ choices for the pair (v_T, c_T) for every $T \in I(B, Z)$, and at most $(s-2)r$ further choices for the pair (v'_T, c'_T) for every $T \in I' \subseteq I(B, Z)$, see Definition 5.3. Moreover, all edges of $\mathcal{T}$ that contain (v, i) must come from one such set B.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Since $V(\mathcal{T}) = Z \times [\![r]\!]$, inequality (6) implies that

$$\begin{aligned}\sum_{v \in V(\mathcal{T})} \deg_{\mathcal{T}}(v)^2 &\leqslant r \cdot \sum_{v \in Z} \gamma(v)^2 \cdot \Delta_1(\mathcal{B})^2 \cdot (rs)^{4L+2K} \\ &\leqslant \frac{K \cdot e(\mathcal{B})^2}{v(\mathcal{H})^2} \cdot (rs)^{4L+2K+1} \cdot \sum_{v \in Z} \gamma(v)^2.\end{aligned} \tag{7}$$

In the remainder of the proof, we will bound the expected value of the sum in the right-hand side of (7).

Fix an arbitrary $v \in V$ and observe that

$$\mathbb{E}[\gamma(v)^2] = \sum_{\substack{A, A' \in \mathcal{H} \\ v \in A \cap A'}} \underbrace{\mathbb{P}\big(|A \cap (Z \setminus \{v\})| = s-2 \text{ and } |A' \cap (Z \setminus \{v\})| = s-2\big)}_{P(A,A')};$$

note that $P(A, A) \leqslant (s-1) \cdot p^{s-2}$ and that, if $A \neq A'$, then

$$P(A, A') \leqslant (s-1)^2 \cdot p^{|A \cup A'|-3} = (s-1)^2 \cdot p^{2s-3-|A \cap A'|}.$$

Consequently,

$$\mathbb{E}[\gamma(v)^2] = \deg_{\mathcal{H}}(v) \cdot (s-1) \cdot p^{s-2} \cdot \left(1 + \sum_{t=1}^{s-1} \binom{s-2}{t-1} \cdot (s-1) \cdot \Delta_t(\mathcal{H}) \cdot p^{s-1-t}\right).$$

Our assumption that $\mathcal{H}$ is nonclustered implies that, for every $t \in [\![s-1]\!]$,

$$\Delta_t(\mathcal{H}) \cdot p^{s-1-t} \leqslant K_{\mathcal{H}}^{s-1} \cdot \Delta_t(\mathcal{H}) \cdot p_{\mathcal{H}}^{s-1-t} \leqslant K_{\mathcal{H}}^{s} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \cdot p_{\mathcal{H}}^{s-2} = \frac{K_{\mathcal{H}}^{s}}{p_{\mathcal{H}}}.$$

and thus

$$\mathbb{E}[\gamma(v)^2] \leqslant \Gamma \cdot \deg_{\mathcal{H}}(v) \cdot p_{\mathcal{H}}^{s-3}$$

for some constant Γ that depends only on $K_{\mathcal{H}}$ and s. Since the variable $\gamma(v)$ is independent of the event $v \in Z$, we have

$$\begin{aligned}\mathbb{E}\left[\sum_{v \in Z} \gamma(v)^2\right] &= \sum_{v \in V} p \cdot \mathbb{E}[\gamma(v)^2] \leqslant \Gamma \cdot K_{\mathcal{H}} \cdot p_{\mathcal{H}}^{s-2} \cdot \sum_{v \in V} \deg_{\mathcal{H}}(v) \\ &= \Gamma \cdot K_{\mathcal{H}} \cdot p_{\mathcal{H}}^{s-2} \cdot s \cdot e(\mathcal{H}) = \frac{\Gamma \cdot K_{\mathcal{H}} \cdot s \cdot v(\mathcal{H})}{p_{\mathcal{H}}}.\end{aligned}$$

Taking expectations of both sides of (7) and substituting the above inequality yields the assertion of the lemma. □

Proof of Lemma 5.8. For every pair of distinct $v, w \in V$, let $\gamma_2(v, w)$ be the number of triples $(A_v, A_w, B) \in \mathcal{H}^2 \times \mathcal{B}$ satisfying the following:

(i) $v \in A_v \setminus B$ and $w \in A_w \setminus B$,
(ii) $|A_v \cap B| = |A_w \cap B| = 1$,
(iii) $A_v \cap A_w \subseteq B$ or $A_v \setminus B = A_w \setminus B$,
(iv) $(A_v \cup A_w) \setminus (B \cup \{v, w\}) \subseteq Z$.

We claim that, for every pair of distinct $v, w \in Z$ and all $i, j \in [\![r]\!]$,

$$\deg_{\mathcal{T}}\{(v, i), (w, j)\} \leqslant \gamma_2(v, w) \cdot (rs)^{2L+K}. \tag{8}$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Indeed, if $v, w \in Z$, then $\gamma_2(v,w)$ is an upper bound on the number of $B \in \mathcal{B}'_Z$ such that $v \in T_v$ and $w \in T_w$ for some $T_v, T_w \in I(B,Z)$. Every such B gives rise to at most $(rs)^{2L+K}$ edges of $\mathcal{T}$, counting multiplicities, and all edges of $\mathcal{T}$ that contain both (v,i) and (w,j) must come from one such set B.

Since $V(\mathcal{T}) = Z \times [\![r]\!]$ and no edge of $\mathcal{T}$ contains a pair of vertices $\{(v,i),(v,j)\}$ with $i \neq j$, inequality (8) implies that

$$\sum_{T \in \binom{V(\mathcal{T})}{2}} \deg_{\mathcal{T}}(T)^2 \leqslant r^2 \cdot \sum_{v,w \in Z} \gamma_2(v,w)^2 \cdot (rs)^{4L+2K}. \tag{9}$$

In the remainder of the proof, we will bound the expected value of the sum in the right-hand side of (9).

Claim 5.9. *For every pair of distinct vertices $v, w \in V$,*

$$\mathbb{E}[\gamma_2(v,w)^2] \ll \frac{e(\mathcal{B})}{p_{\mathcal{H}} \cdot v(\mathcal{H})} \cdot \mathbb{E}[\gamma_2(v,w)].$$

Proof. Let $\mathcal{X}$ denote the family of all triples (A_v, A_w, B) that satisfy i–iii above and observe that

$$\gamma_2(v,w) = \sum_{(A_v,A_w,B) \in \mathcal{X}} \mathbb{1}_{(A_v \cup A_w) \setminus (B \cup \{v,w\}) \subseteq Z},$$

see (iv). Therefore, it suffices to show that, for each $(A_v, A_w, B) \in \mathcal{X}$,

$$\mathbb{E}\big[\gamma_2(v,w) \mid (A_v \cup A_w) \setminus (B \cup \{v,w\}) \subseteq Z\big] \ll \frac{\Delta_1(\mathcal{B})}{p_{\mathcal{H}}}. \tag{10}$$

We partition the family $\mathcal{X}$ according to the intersection pattern with our chosen triple (A_v, A_w, B). First, for every $t \in \{0, \dots, s-3\}$, denote by $\mathcal{X}_t$ the family of all triples (A'_v, A'_w, B') such that $A'_v \setminus B = A'_w \setminus B$ and A'_v intersects $(A_v \cup A_w) \setminus (B \cup \{v,w\})$ in exactly t elements. Second, for every $(t_v, t_w) \in \{0, \dots, s-2\}^2$, denote by $\mathcal{X}_{t_v,t_w}$ the family of all triples (A'_v, A'_w, B') such that $A'_v \cap A'_w \subseteq B$ and, for $u \in \{v,w\}$, the set A'_u intersects $(A_v \cup A_w) \setminus (B \cup \{v,w\})$ in exactly t_u elements. Observe that

$$\mathcal{X} = \bigcup_{t=0}^{s-3} \mathcal{X}_t \cup \bigcup_{t_v,t_w=0}^{s-2} \mathcal{X}_{t_v,t_w}.$$

Further, note that, for every $t \in \{0, \dots, s-3\}$,

$$|\mathcal{X}_t| \leqslant \Delta_{t+2}(\mathcal{H}) \cdot \Delta_1(\mathcal{B}) \cdot \max_{B \in \mathcal{B}} |B| \ll p_{\mathcal{H}}^{t+1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \cdot \Delta_1(\mathcal{B}) = p_{\mathcal{H}}^{t+2-s} \cdot \Delta_1(\mathcal{B}).$$

Crucially, we claim that, for every $t_v, t_w \in \{0, \dots, s-2\}$,

$$\begin{aligned} |\mathcal{X}_{t_v,t_w}| &\leqslant \min\big\{\Delta_{t_v+1}(\mathcal{H}) \cdot \Delta_{t_w+2}(\mathcal{H}), \Delta_{t_w+1}(\mathcal{H}) \cdot \Delta_{t_v+2}(\mathcal{H})\big\} \cdot \Delta_1(\mathcal{B}) \\ &\ll p_{\mathcal{H}}^{t_v+t_w+1} \cdot \left(\frac{e(\mathcal{H})}{v(\mathcal{H})}\right)^2 \cdot \Delta_1(\mathcal{B}) = p_{\mathcal{H}}^{t_v+t_w-2s+3} \cdot \Delta_1(\mathcal{B}). \end{aligned}$$

Indeed, $\Delta_t(\mathcal{H}) \leqslant K p_{\mathcal{H}}^{t-1} \cdot e(\mathcal{H})/v(\mathcal{H})$ for all $t \in [\![s]\!]$ and $\Delta_t(\mathcal{H}) \ll p_{\mathcal{H}}^{t-1} \cdot e(\mathcal{H})/v(\mathcal{H})$ when $2 \leqslant t \leqslant s-1$ and one of $\max\{t_v, t_w\}+1$ or $\min\{t_v, t_2\}+2$ must belong to $\{2, \dots, s-1\}$. Since $p = O(p_{\mathcal{H}})$, estimate (10) follows after noting that, first, for every $t \in \{0, \dots, s-3\}$ and every $(A'_v, A'_w, B') \in \mathcal{X}_t$,

$$\mathbb{P}\big((A'_v \cup A'_w) \setminus (B' \cup \{v,w\}) \subseteq Z \mid (A_v \cup A_w) \setminus (B \cup \{v,w\}) \subseteq Z\big) = p^{s-3-t},$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

and, second, for all $(t_v, t_w) \in \{0, \dots, s-2\}^2$ and every $(A'_v, A'_w, B') \in \mathcal{X}_{t_v,t_w}$,

$$\mathbb{P}\big((A'_v \cup A'_w) \setminus (B' \cup \{v,w\}) \subseteq Z \mid (A_v \cup A_w) \setminus (B \cup \{v,w\}) \subseteq Z\big) = p^{2s-4-t_v-t_w}.$$

Since $p \leqslant K_{\mathcal{H}} \cdot p_{\mathcal{H}}$ and $\Delta_1(\mathcal{B}) \leqslant K \cdot e(\mathcal{B})/v(\mathcal{H})$, this concludes the proof of the claim. □

Claim 5.10. *We have*

$$\mathbb{E}\left[\sum_{v,w\in Z} \gamma_2(v,w)\right] = O\big(e(\mathcal{B})\big).$$

Proof. Let $\mathcal{X}_1$ denote the family of all pairs $(A,B) \in \mathcal{H} \times \mathcal{B}$ such that $|A \cap B| = 1$ and let $\mathcal{X}_2$ denote the family of all triples $(A', A'', B') \in \mathcal{H}^2 \times \mathcal{B}$ such that $|A' \cap B'| = |A'' \cap B'| = 1$ and $A' \cap A'' \subseteq B'$. Note that

$$\sum_{v,w\in Z} \gamma_2(v,w) = \binom{s-1}{2} \cdot \Big|\big\{(A,B) \in \mathcal{X}_1 : A \setminus B \subseteq Z\big\}\Big| + (s-1)^2 \cdot \Big|\big\{(A', A'', B') \in \mathcal{X} : (A' \cup A'') \setminus B' \subseteq Z\big\}\Big|. \tag{11}$$

The claim now follows as

$$|\mathcal{X}_1| \leqslant \sum_{B\in\mathcal{B}} |B| \cdot \Delta_1(\mathcal{H}) = O\left(e(\mathcal{B}) \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}\right),$$
$$|\mathcal{X}_2| \leqslant \sum_{B\in\mathcal{B}} |B|^2 \cdot \Delta_1(\mathcal{H})^2 = O\left(e(\mathcal{B}) \cdot \frac{e(\mathcal{H})^2}{v(\mathcal{H})^2}\right),$$

and, for all $(A,B) \in \mathcal{X}_1$ and $(A', A'', B') \in \mathcal{X}_2$,

$$\mathbb{P}\big(A \setminus B \subseteq Z\big) = p^{s-1} = O\big(p_{\mathcal{H}}^{s-1}\big),$$
$$\mathbb{P}\big((A' \cup A'') \setminus B' \subseteq Z\big) = p^{2s-2} = O\big(p_{\mathcal{H}}^{2s-2}\big).$$

Indeed, since $p_{\mathcal{H}}^{s-1} \cdot e(\mathcal{H}) = v(\mathcal{H})$, taking expectations of both sides of (11) and substituting the above estimates yields the claimed estimate. □

Since the variable $\gamma_2(v,w)$ is independent of the event $\{v,w\} \subseteq Z$, we have

$$\mathbb{E}\left[\sum_{v,w\in Z} \gamma_2(v,w)^2\right] = p^2 \cdot \sum_{v,w\in V} \mathbb{E}[\gamma_2(v,w)^2] \ll p^2 \cdot \frac{e(\mathcal{B})}{p_{\mathcal{H}} \cdot v(\mathcal{H})} \cdot \sum_{v,w\in V} \mathbb{E}[\gamma_2(v,w)]$$
$$= \frac{e(\mathcal{B})}{p_{\mathcal{H}} \cdot v(\mathcal{H})} \cdot \mathbb{E}\left[\sum_{v,w\in Z} \gamma_2(v,w)\right] = O\left(\frac{e(\mathcal{B})^2}{p_{\mathcal{H}} \cdot v(\mathcal{H})}\right),$$

where the first inequality follows from Claim 5.9 and the second inequality follows from Claim 5.10. The assertion of the lemma follows by taking expectations of both sides of (9) and substituting the above inequality. □

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

5.6. Containers for colourings

Let σ be the sequence from the statement of Lemma 5.8 and let $\tau = \sigma^{1/(2L+1)} = o(1)$. Let $\mathcal{Z}''$ be the collection of all $Z \in \mathcal{Z}'$ such that

$$\sum_{v \in V(\mathcal{T})} \deg_{\mathcal{T}}(v)^2 \leqslant \Gamma_{\mathcal{H}} \cdot \frac{e(\mathcal{B})^2}{|Z|} \qquad \text{and} \qquad \sum_{T \in \binom{V(\mathcal{T})}{2}} \deg_{\mathcal{T}}(T)^2 \leqslant \tau^{2L} \cdot \frac{e(\mathcal{B})^2}{|Z|}.$$

Since $\mathcal{Z}'' \subseteq \mathcal{Z}$, the set Z satisfies (i) in the statement of Theorem 5.

Lemma 5.11. *Every $Z \in \mathcal{Z}''$ satisfies (ii) in the statement of Theorem 5.*

Before we prove the lemma, let us point out that it implies the assertion of the theorem. Indeed, it follows from Lemmas 5.7 and 5.8, Markov's inequality, and standard estimates for the tails of the binomial distribution that, when $Z \sim V_p$,

$$\begin{aligned}\mathbb{P}(Z \in \mathcal{Z}'') &\geqslant \mathbb{P}(Z \in \mathcal{Z}') - \mathbb{P}\big(|Z| \geqslant 2p_{\mathcal{H}} v(\mathcal{H})\big) - 2\Gamma/\Gamma_{\mathcal{H}} - 2\sigma \cdot \tau^{-2L} \\ &\geqslant \varepsilon/2 - o(1) - 2\Gamma/\Gamma_{\mathcal{H}} - 2\tau \geqslant \varepsilon/4.\end{aligned}$$

Proof of Lemma 5.11. Suppose that $Z \in \mathcal{Z}''$. We will construct the desired family of partial colourings of Z by applying the container lemma, Theorem B, to a collection of uniform subhypergraphs of the hypergraph $\mathcal{T}$.

To this end, note first that Proposition 5.6, invoked twice, implies that $\mathcal{T}$ contains a subhypergraph $\mathcal{T}'$ with at least $e(\mathcal{T}) - 2 \cdot (\varepsilon/16) \cdot e(\mathcal{B})$ edges that satisfies

$$\Delta_1(\mathcal{T}') \leqslant \frac{16\Gamma_{\mathcal{H}}}{\varepsilon} \cdot \frac{e(\mathcal{B})}{|Z|} \qquad \text{and} \qquad \Delta_2(\mathcal{T}') \leqslant \frac{16\tau^{2L}}{\varepsilon} \cdot \frac{e(\mathcal{B})}{|Z|} \leqslant \tau^{2L-1} \cdot \frac{e(\mathcal{B})}{|Z|}.$$

Fix one such hypergraph $\mathcal{T}'$ and note that

$$e(\mathcal{T}') \leqslant r \cdot |Z| \cdot \Delta_1(\mathcal{T}') \leqslant \frac{16\Gamma_{\mathcal{H}} r}{\varepsilon} \cdot e(\mathcal{B}). \tag{12}$$

Let $\mathcal{T}_1, \dots, \mathcal{T}_{2L}$ be the subhypergraphs of $\mathcal{T}'$ that comprise all edges of $\mathcal{T}'$ of cardinalities $1, \dots, 2L$, respectively, and note that $\mathcal{T}' = \mathcal{T}_1 \cup \cdots \cup \mathcal{T}_{2L}$, since every edge of $\mathcal{T}$ has cardinality at most

$$2 \cdot \max_{B \in \mathcal{B}'_Z} |I(B, Z)| \leqslant 2L,$$

see Definition 5.3. Define

$$U \coloneqq \big\{u \in [\![2L]\!] : e(\mathcal{T}_u) \geqslant \varepsilon/(16L) \cdot e(\mathcal{B})\big\}$$

and fix an arbitrary $u \in U$. Since

$$\Delta_1(\mathcal{T}_u) \leqslant \Delta_1(\mathcal{T}') \leqslant \frac{16\Gamma_{\mathcal{H}}}{\varepsilon} \cdot \frac{e(\mathcal{B})}{|Z|} \leqslant \frac{256\Gamma_{\mathcal{H}} Lr}{\varepsilon^2} \cdot \frac{e(\mathcal{T}_u)}{v(\mathcal{T}_u)}$$

and, for every $\ell \in \{2, \dots, u\}$,

$$\Delta_\ell(\mathcal{T}_u) \leqslant \Delta_2(\mathcal{T}') \leqslant \tau^{2L-1} \cdot \frac{e(\mathcal{B})}{|Z|} \leqslant \frac{16Lr}{\varepsilon} \cdot \tau^{\ell-1} \cdot \frac{e(\mathcal{T}_u)}{v(\mathcal{T}_u)},$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

we may apply Theorem B, with $k_{4.7} \coloneqq u$,

$$K_{4.7} \coloneqq \max\left\{\frac{256\Gamma_{\mathcal{H}}Lr}{\varepsilon^2}, \frac{16Lr}{\varepsilon}\right\}, \quad \text{and} \quad \varepsilon_{4.7} \coloneqq \frac{\varepsilon^2}{256\Gamma_{\mathcal{H}}Lr} \overset{(12)}{\leqslant} \frac{\varepsilon/(16L)\cdot e(\mathcal{B})}{e(\mathcal{T}')}$$

to get an integer t (that depends only on ε, $\Gamma_{\mathcal{H}}$, r, and L) and a collection $\mathcal{C}_u$ of at most

$$\left(\sum_{i=0}^{\tau r|Z|}\binom{r|Z|}{i}\right)^t \leqslant \left(\frac{er|Z|}{\tau r|Z|}\right)^{t\tau r|Z|} = \exp\big(o(|Z|)\big)$$

subsets of $Z \times [\![r]\!]$ with the following properties:

(i) Every proper colouring $\psi\colon Z \to [\![r]\!]$ of $\mathcal{H}[Z]$, viewed as a subset of $Z \times [\![r]\!]$, is contained in a member of $\mathcal{C}_u$.
(ii) Every member of $\mathcal{C}_u$ induces fewer than $\varepsilon/(16L)\cdot e(\mathcal{B})$ edges in $\mathcal{T}_u$.

Finally, we let Ψ to be the collection of all partial colourings ψ_C, where C is a set of the form

$$C = \bigcap_{u\in U} C_u,$$

where, for each $u \in U$, the set C_u is a member of $\mathcal{C}_u$, such that $\pi_1(C) = Z$; note that

$$|\Psi| \leqslant \prod_{u\in U}|\mathcal{C}_u| = \exp\big(o(|Z|)\big).$$

It follows from (i) that, for each $u \in U$, every proper colouring $\psi\colon Z \to [\![r]\!]$ of $\mathcal{H}[Z]$ is contained in some set $C_u \in \mathcal{C}_u$ and, consequently, in some set C of the above form; in particular ψ extends the partial colouring $\psi_C \in \Psi$. We now show that, for every partial colouring $\psi_C \in \Psi$, we have $|F(\psi_C)| = \Omega\big(v(\mathcal{H})\big)$; this will conclude the proof of the lemma, as every vertex in $F(\psi_C)$ is the centre of a rainbow star.

To this end, choose an arbitrary set C of the above form. Since $e(\mathcal{T}_u[C]) \leqslant e(\mathcal{T}_u[C_u]) \leqslant \varepsilon/(16L)\cdot e(\mathcal{B})$ when $u \in U$, see (ii) above, and $e(\mathcal{T}_u[C]) \leqslant e(\mathcal{T}_u) < \varepsilon/(16L)\cdot e(\mathcal{B})$ when $u \notin U$, we have

$$e(\mathcal{T}'[C]) = \sum_{u=1}^{2L} e(\mathcal{T}_u[C]) \leqslant 2L\cdot\varepsilon/(16L)\cdot e(\mathcal{B}) = (\varepsilon/8)\cdot e(\mathcal{B}).$$

Consequently,

$$e(\mathcal{T}[C]) \leqslant e(\mathcal{T}'[C]) + e(\mathcal{T}) - e(\mathcal{T}') \leqslant (\varepsilon/4)\cdot e(\mathcal{B})$$

and thus Lemma 5.5 implies that

$$|F(\psi_C)| \geqslant \frac{e(\mathcal{B}'_Z) - e(\mathcal{T}[C])}{\Delta_1(\mathcal{B})} \geqslant \frac{\varepsilon\cdot e(\mathcal{B})}{4\cdot\Delta_1(\mathcal{B})} \geqslant \frac{\varepsilon}{4K}\cdot v(\mathcal{H}),$$

since $Z \in \mathcal{Z}'$ implies that $e(\mathcal{B}'_Z) \geqslant (\varepsilon/2)\cdot e(\mathcal{B})$. □

6. Rainbow stars and constellations

In this section, we prove Theorem 3.8, which encapsulates Step (III) from the proof outline. We restate the theorem here for the reader's convenience.

Theorem 3.8. *Suppose that $s \geqslant 3$, $r \geqslant 2$, and $\varepsilon > 0$. Let $\mathcal{H}$ be a sequence of s-uniform hypergraphs that satisfies assumptions (A2) and (A5) and suppose that $Z \sim V(\mathcal{H})_p$ for some $p = \Theta(p_{\mathcal{H}})$. With*

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

probability at least $1-\varepsilon$*, for every partial colouring* ψ *of* Z *with elements of* $[\![r]\!]$*, if* $\Omega\big(v(\mathcal{H})\big)$ *vertices of* $\mathcal{H}$ *are centres of rainbow stars, then* $\Omega\big(e(\mathcal{H})\big)$ *edges of* $\mathcal{H}$ *are bases of rainbow constellations.*

The heart of the proof will be transferring the rainbow star-constellation property from $\mathcal{H}$ to $\mathcal{H}_p$. We will make use of the container lemma (Theorem B) to prove the following sparse random analogue of the rainbow star-constellation property.

Theorem 6.1. *Suppose that* $r \geqslant 2$*,* $s \geqslant 3$*, and* $\mathcal{H}$ *is a sequence of* s*-uniform hypergraphs that satisfies assumptions (A2) and (A5). For every positive* β_s*, there exists a positive* β_c *such that the following holds: Suppose that* $p = \Omega(p_{\mathcal{H}})$ *and let* $Z \sim V(\mathcal{H})_p$*. With probability* $1-o(1)$*, every partial* $[\![r]\!]$*-colouring of* Z *that admits at least* $\beta_s \cdot (p/p_{\mathcal{H}})^{(r-1)(s-1)} v(\mathcal{H})$ *rainbow stars must also admit at least* $\beta_c \cdot (p/p_{\mathcal{H}})^{(r-1)(s-1)s} e(\mathcal{H})$ *rainbow constellations.*

The above theorem will be combined with the following technical lemma, which implies that, in a typical set $Z \sim V(\mathcal{H})_p$, a family of $\Omega\big((p/p_{\mathcal{H}})^{(r-1)(s-1)s} e(\mathcal{H})\big)$ rainbow constellations has to determine $\Omega(e(\mathcal{H}))$ different base edges. We will say that a set $Z \subseteq V(\mathcal{H})$ *admits* a (noncoloured) star $\{A_1, \dots, A_k\}$ with centre v if $(A_1 \cup \cdots \cup A_k) \setminus \{v\} \subseteq Z$; a set Z admits a constellation if it admits each of the s stars that comprise it. Given a set $Z \subseteq V(\mathcal{H})$ and an edge $A \in \mathcal{H}$, we denote by $\mathrm{con}(A, Z)$ the number of constellations with base edge A that are admitted by Z.

Lemma 6.2. *Suppose that* $r \geqslant 2$*,* $s \geqslant 3$*, and* $\mathcal{H}$ *is a sequence of* s*-uniform hypergraphs that satisfies assumption (A2). For every* $c > 0$*, there exists a constant* Γ *such that the following holds: If* $p \geqslant c \cdot p_{\mathcal{H}}$ *and* $Z \sim V(\mathcal{H})_p$*, then, for every* $A \in \mathcal{H}$*,*

$$\mathbb{E}\left[\sum_{A \in \mathcal{H}} \mathrm{con}(A, Z)^2\right] \leqslant \Gamma \cdot \left(\frac{p}{p_{\mathcal{H}}}\right)^{2(r-1)(s-1)s} \cdot e(\mathcal{H}).$$

We postpone the proofs of Theorem 6.1 and Lemma 6.2 to later subsections and first show how they imply Theorem 3.8.

Proof of Theorem 3.8. Suppose that $s \geqslant 3$, $r \geqslant 2$, and $\varepsilon > 0$ and let $\mathcal{H}$ be a sequence of nonclustered, s-uniform hypergraphs that satisfies the star-constellation property with r colours. Assume that $p = \Theta(p_{\mathcal{H}})$ and $Z \sim V(\mathcal{H})_p$. It follows from Lemma 6.2 and Markov's inequality that, for some constant C,

$$\sum_{A \in \mathcal{H}} \mathrm{con}(A, Z)^2 \leqslant C e(\mathcal{H})$$

with probability at least $1-\varepsilon/2$. By Theorem 6.1, as $p/p_{\mathcal{H}} = \Theta(1)$, with probability at least $1-\varepsilon/2$, every partial $[\![r]\!]$-colouring of Z that admits $\Omega\big(v(\mathcal{H})\big)$ rainbow stars must also admit $\Omega\big(e(\mathcal{H})\big)$ rainbow constellations. Assume that the random set Z has both these properties, which happens with probability at least $1-\varepsilon$. Fix any colouring ψ with $\Omega\big(v(\mathcal{H})\big)$ rainbow stars and let $\mathcal{H}_\psi$ comprise all edges of $\mathcal{H}$ that are the base of at least one rainbow constellation. On the one hand, since every $A \in \mathcal{H}_\psi$ is the base of at most $\mathrm{con}(A, Z)$ rainbow constellations, our assumption on Z implies that

$$\sum_{A \in \mathcal{H}_\psi} \mathrm{con}(A, Z) \geqslant \Omega\big(e(\mathcal{H})\big).$$

On the other hand, the Cauchy–Schwarz inequality yields

$$\left(\sum_{A \in \mathcal{H}_\psi} \mathrm{con}(A, Z)\right)^2 \leqslant e(\mathcal{H}_\psi) \cdot \sum_{A \in \mathcal{H}_\psi} \mathrm{con}(A, Z)^2 \leqslant e(\mathcal{H}_\psi) \cdot \sum_{A \in \mathcal{H}} \mathrm{con}(A, Z)^2.$$

Combining the three displayed inequalities gives $e(\mathcal{H}_\psi) = \Omega\big(e(\mathcal{H})\big)$, as desired. □

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

6.1. The hypergraphs of stars and constellations

Assume from now on that $s \geqslant 3$ and $r \geqslant 2$ are integers and that $\mathcal{H}$ is a sequence of s-uniform hypergraphs that satisfies assumptions (A2) and (A5). In order to use the container lemma in the context of Theorem 6.1, let us define two (multi)hypergraphs with vertex set $V(\mathcal{H}) \times [\![r]\!]$:

- the $(r-1)(s-1)$-uniform hypergraph $\mathcal{R}_\star$ of rainbow stars,
- the $(r-1)(s-1)s$-uniform hypergraph $\mathcal{R}_\star^\star$ of rainbow constellations.

The edges of $\mathcal{R}_\star$ are all sets of the form

$$E_i\Big((A_j)_{j\in[\![r]\!]\setminus\{i\}}\Big) = \bigcup_{j\in[\![r]\!]\setminus\{i\}} \big(A_j \setminus \{v\}\big) \times \{j\},$$

where $i \in [\![r]\!]$ and $\{A_j : j \in [\![r]\!] \setminus \{i\}\}$ is an $(r-1)$-star with centre v. We add such sets to the multihypergraph $\mathcal{R}_\star$ with their proper multiplicities, that is, $E_i(A_1, \dots, A_{r-1})$ is added to $\mathcal{R}_\star$ once for every ordering of the edges of each $(r-1)$-star $\{A_1, \dots, A_{r-1}\}$.[8] The edges of $\mathcal{R}_\star^\star$ are all sets of the form

$$E_i(\mathbf{A}_1) \cup \cdots \cup E_i(\mathbf{A}_s),$$

where $\mathbf{A}_1, \dots, \mathbf{A}_s$ are arbitrary orderings of the edge sets of some s stars that form a constellation. As in the case of $\mathcal{R}_\star$, every edge of $\mathcal{R}_\star^\star$ appears with its proper multiplicity, that is, once for every choice of a base edge formed by the centres of the stars in the constellation.

Our next two technical lemmas provide lower bounds on the numbers of edges in the hypergraphs $\mathcal{R}_\star$ and $\mathcal{R}_\star^\star$ and upper bounds on the sequence of maximum degrees of the latter hypergraph. These estimates will allow us to apply the container lemma (Theorem B) to the $(r-1)(s-1)s$-uniform hypergraph $\mathcal{R}_\star^\star$ and construct a family of only $\exp\big(o(p \cdot v(\mathcal{H}))\big)$ containers that cover the family of all partial colourings of $V(\mathcal{H})$ that admit many fewer than $p^{(r-1)(s-1)s} \cdot e(\mathcal{R}_\star^\star)$ rainbow constellations.

Lemma 6.3. *There is a positive constant c that depends only on r and s such that*

$$e(\mathcal{R}_\star) \geqslant c \cdot \frac{e(\mathcal{H})^{r-1}}{v(\mathcal{H})^{r-2}} \qquad \text{and} \qquad e(\mathcal{R}_\star^\star) \geqslant c \cdot \frac{e(\mathcal{H})^{(r-1)s+1}}{v(\mathcal{H})^{(r-1)s}}.$$

Lemma 6.4. *There are constants Γ and Γ' that depend only on r and the sequence $\mathcal{H}$ and such that*

$$\Delta_1(\mathcal{R}_\star^\star) \leqslant \Gamma' \cdot \left(\frac{e(\mathcal{H})}{v(\mathcal{H})}\right)^{(r-1)s+1} \leqslant \Gamma \cdot \frac{e(\mathcal{R}_\star^\star)}{v(\mathcal{R}_\star^\star)}.$$

Moreover, for every $t \geqslant 2$,

$$\Delta_t(\mathcal{R}_\star^\star) \ll p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{R}_\star^\star)}{v(\mathcal{R}_\star^\star)}.$$

The next lemma translates the rainbow star-constellation property into the language of $\mathcal{R}_\star$ and $\mathcal{R}_\star^\star$.

Lemma 6.5. *For every positive constant β, there exists a positive constant ε that depends only on β, r, and the sequence $\mathcal{H}$ and such that the following holds for every $C \subseteq V(\mathcal{H}) \times [\![r]\!]$: If $e\big(\mathcal{R}_\star[C]\big) \geqslant \beta \cdot e(\mathcal{R}_\star)$, then $e\big(\mathcal{R}_\star^\star[C]\big) \geqslant \varepsilon \cdot e(\mathcal{R}_\star^\star)$.*

Furthermore, it will be convenient to define the $(r-1)(s-1)$-uniform (multi)hypergraph $\mathcal{M}_\star$ of uncoloured $(r-1)$-stars whose vertex set is $V(\mathcal{H})$. The edges of $\mathcal{M}_\star$ are all sets of the form

$$(A_1 \cup \cdots \cup A_{r-1}) \setminus \{v\},$$

[8] It is possible that $E_i(A_1, \dots, A_{r-1}) = E_i(A'_1, \dots, A'_{r-1})$ when $\{A_1, \dots, A_{r-1}\}$ and $\{A'_1, \dots, A'_{r-1}\}$ are two different stars that differ only in their centres.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

where $A_1, \dots, A_{r-1}$ form a star in $\mathcal{H}$ whose centre vertex is v. Observe that $\mathcal{M}_\star$ can be thought of as the image of $\mathcal{R}_\star$ via the projection $\pi_1 \colon V(\mathcal{H}) \times [\![r]\!] \to V(\mathcal{H})$ of its vertex set onto the first coordinate. More precisely, since every star admits $r!$ colourings that make it rainbow, we have that $\pi_1(\mathcal{R}_\star) = r! \cdot \mathcal{M}_\star$ as multihypergraphs. This allows us to bound degrees of subhypergraphs of $\mathcal{R}_\star$ induced by various $[\![r]\!]$-colourings of Z from above using respective degrees of $\mathcal{M}_\star[Z]$. The advantage of such an approach is that the latter degrees depend only on Z rather than on the particular colouring of Z. This motivates the two final lemmas of this subsection.

Lemma 6.6. *There is a constant Γ that depends only on r and the sequence $\mathcal{H}$ such that*

$$\Delta_1(\mathcal{M}_\star) \leqslant \Gamma \cdot \frac{e(\mathcal{M}_\star)}{v(\mathcal{M}_\star)} \qquad \text{and} \qquad e(\mathcal{M}_\star) \leqslant \Gamma \cdot \frac{e(\mathcal{H})^{r-1}}{v(\mathcal{H})^{r-2}}.$$

Moreover, for every $t \geqslant 2$,

$$\Delta_t(\mathcal{M}_\star) \ll p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{M}_\star)}{v(\mathcal{M}_\star)}.$$

Recall the definition of pseudo-variance (Definition 4.8) given in Section 4.3. In our final lemma, we abuse the notation somewhat and write $\mathrm{Var}'(\mathcal{M}_\star[Z])$ to denote the pseudo-variance of the sequence of events $(\{A \subseteq Z\})_{A \in \mathcal{M}_\star}$.

Lemma 6.7. *For every positive constant c, there is a constant Γ that depends only on c, r, and the sequence $\mathcal{H}$ such that the following holds: If $Z \sim V(\mathcal{H})_p$ for some $p \geqslant c \cdot p_{\mathcal{H}}$, then*

$$\mathrm{Var}'\big(\mathcal{M}_\star[Z]\big) \leqslant \Gamma \cdot \frac{\mathbb{E}[e(\mathcal{M}_\star[Z])]^2}{p v(\mathcal{M}_\star)}.$$

6.2. Proofs of Lemmas 6.2–6.7

In this subsection, we prove Lemmas 6.2–6.7. All proofs are straightforward, albeit somewhat technical. Since the proofs of Lemmas 6.4 and 6.6 are very similar, but the latter is simpler, we present them in reverse order. Throughout this section, we assume that $r \geqslant 2$ and $s \geqslant 3$ are integers and that $\mathcal{H}$ is a sequence of s-uniform hypergraphs that satisfies assumptions (A2) and (A5). In particular, there is a constant K such that, for all $t \geqslant 1$,

$$\Delta_t(\mathcal{H}) \leqslant K \cdot p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}. \tag{13}$$

Proof of Lemma 6.2. Fix an edge $A \in \mathcal{H}$ and let $\mathcal{X}_A$ be the family of all constellations with base edge A. Denoting by $\mathrm{supp}(\mathcal{C})$ the union of all $(r-1)s$ edges forming a constellation $\mathcal{C}$, we have

$$\begin{aligned}\mathbb{E}\big[\mathrm{con}(A,Z)^2\big] &= \sum_{\mathcal{C},\mathcal{C}' \in \mathcal{X}_A} p^{|(\mathrm{supp}(\mathcal{C}) \cup \mathrm{supp}(\mathcal{C}')) \setminus A|} \\ &= \sum_{\mathcal{C} \in \mathcal{X}_A} p^{|\mathrm{supp}(\mathcal{C}) \setminus A|} \cdot \sum_{\mathcal{C}' \in \mathcal{X}_A} p^{|\mathrm{supp}(\mathcal{C}') \setminus \mathrm{supp}(\mathcal{C})|}.\end{aligned}$$

For every $\mathcal{C} \in \mathcal{X}_A$ and $t \in \{0, \dots, (r-1)(s-1)s\}$, let $\mathcal{X}_{\mathcal{C},t}$ denote the family of all $\mathcal{C}' \in \mathcal{X}_A$ satisfying $|\mathrm{supp}(\mathcal{C}') \cap \mathrm{supp}(\mathcal{C})| = |A| + t$, so that

$$\mathcal{X}_A = \bigcup_{t=0}^{(r-1)(s-1)s} \mathcal{X}_{\mathcal{C},t}$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

and, consequently,

$$\mathbb{E}\big[\mathrm{con}(A,Z)^2\big] = \sum_{\mathcal{C}\in\mathcal{X}_A} p^{(r-1)(s-1)s} \cdot \sum_{t=0}^{(r-1)(s-1)s} |\mathcal{X}_{\mathcal{C},t}| \cdot p^{(r-1)(s-1)s-t}.$$

We may further partition each $\mathcal{X}_{\mathcal{C},t}$ according to the intersection pattern of $\mathcal{C}' \in \mathcal{X}_{\mathcal{C},t}$ with $\mathrm{supp}(\mathcal{C})$. Namely, for every sequence $\mathbf{t} = (t_{i,j} : i \in [\![r-1]\!], j \in [\![s]\!])$ with $0 \leqslant t_{i,j} \leqslant s-1$, we let $\mathcal{X}_{\mathcal{C},\mathbf{t}}$ be the family of all $\mathcal{C}' \in \mathcal{X}_A$ such that the ith edge of the jth star in $\mathcal{C}'$ (in some arbitrary ordering) intersects $\mathrm{supp}(\mathcal{C}) \setminus A$ in $t_{i,j}$ vertices, so that

$$\mathcal{X}_{\mathcal{C},t} = \bigcup_{\mathbf{t}:\sum\mathbf{t}=t} \mathcal{X}_{\mathcal{C},\mathbf{t}}.$$

Since each edge of every star comprising each $\mathcal{C}' \in \mathcal{X}_A$ intersects A in one vertex, there is a constant Γ' that depends only on r and s such that

$$\begin{aligned} |\mathcal{X}_{\mathcal{C},\mathbf{t}}| &\leqslant \Gamma' \cdot \prod_{i,j} \Delta_{t_{i,j}+1}(\mathcal{H}) \leqslant \Gamma' \cdot p_{\mathcal{H}}^{\sum_{i,j} t_{i,j}} \cdot \left(K \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}\right)^{(r-1)s} \\ &= \Gamma' \cdot K^{(r-1)s} \cdot p_{\mathcal{H}}^{\sum\mathbf{t}-(r-1)(s-1)s}. \end{aligned}$$

In particular, there is a constant Γ'' that depends only on r, s, and K such that

$$\begin{aligned} \mathbb{E}\big[\mathrm{con}(A,Z)^2\big] &\leqslant \Gamma'' \cdot \sum_{\mathcal{C}\in\mathcal{X}_A} p^{(r-1)(s-1)s} \cdot \sum_{t=0}^{(r-1)(s-1)s} \left(\frac{p}{p_{\mathcal{H}}}\right)^{(r-1)(s-1)s-t} \\ &\leqslant \Gamma'' \cdot |\mathcal{X}_A| \cdot \left(\frac{p^2}{p_{\mathcal{H}}}\right)^{(r-1)(s-1)s} \cdot \sum_{t=0}^{(r-1)(s-1)s} c^{-t}. \end{aligned} \tag{14}$$

Since

$$|\mathcal{X}_A| \leqslant \Delta_1(\mathcal{H})^{(r-1)s} \leqslant \left(K \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}\right)^{(r-1)s} = K^{(r-1)s} \cdot p_{\mathcal{H}}^{-(r-1)(s-1)s},$$

summing (14) over all $A \in \mathcal{H}$ yields

$$\mathbb{E}\left[\sum_{A\in\mathcal{H}} \mathrm{con}(A,Z)^2\right] \leqslant \Gamma \cdot \left(\frac{p}{p_{\mathcal{H}}}\right)^{2(r-1)(s-1)s} \cdot e(\mathcal{H}),$$

for some constant Γ that depends only on c, r, s, and K, as claimed. □

Proof of Lemma 6.3. Let $\mathcal{H}'$ be the hypergraph obtained from $\mathcal{H}$ by iteratively removing vertices with degree smaller than $e(\mathcal{H})/(2v(\mathcal{H}))$. Observe that

$$e(\mathcal{H}') > e(\mathcal{H}) - v(\mathcal{H}) \cdot \frac{e(\mathcal{H})}{2v(\mathcal{H})} = \frac{e(\mathcal{H})}{2}$$

and, consequently, that $\delta(\mathcal{H}') \geqslant e(\mathcal{H})/(2v(\mathcal{H}))$. In particular, there are at least

$$e(\mathcal{H}') \cdot \delta(\mathcal{H}')^{r-2} \geqslant 2^{-r} \cdot \frac{e(\mathcal{H})^{r-1}}{v(\mathcal{H})^{r-2}}$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

sequences $A_1, \dots, A_{r-1}$ of edges of $\mathcal{H}'$ that satisfy $A_1 \cap \dots \cap A_{r-1} \neq \emptyset$. If $r = 2$, then each such sequence corresponds to $2s$ edges of $\mathcal{R}_\star$ (there are two colours and s different choices for the centre vertex that makes an edge into a star). Otherwise, if $r \geqslant 3$, then all but at most

$$O(1) \cdot e(\mathcal{H}) \cdot \Delta_1(\mathcal{H})^{r-3} \cdot \Delta_2(\mathcal{H}) \ll \frac{e(\mathcal{H})^{r-1}}{v(\mathcal{H})^{r-2}}$$

of those sequences are orderings of the edges of an $(r-1)$star (whose centre is the unique element of $A_1 \cap \dots \cap A_{r-1}$). Moreover, as $A_1, \dots, A_{r-1}$ range over all such sequences and i ranges over $[\![r]\!]$, the sets $E_i(A_1, \dots, A_{r-1})$ are distinct edges of $\mathcal{R}_\star$.

Similarly, there are at least

$$e(\mathcal{H}') \cdot \delta(\mathcal{H}')^{(r-1)s} \geqslant 2^{-rs} \cdot \frac{e(\mathcal{H})^{(r-1)s+1}}{v(\mathcal{H})^{(r-1)s}}$$

pairs comprising an edge $A = \{v_1, \dots, v_s\}$ of $\mathcal{H}'$ and a set $\{\mathbf{A}_1, \dots, \mathbf{A}_s\}$ of sequences of $r-1$ edges of $\mathcal{H}'$ such that, letting $\mathbf{A}_j = (A_{j,1}, \dots, A_{j,r-1})$, we have $v_j \in A_{j,1} \cap \dots \cap A_{j,r-1}$. Moreover, for all but at most

$$O(1) \cdot e(\mathcal{H}) \cdot \Delta_1(\mathcal{H})^{s(r-1)-1} \cdot \Delta_2(\mathcal{H}) \ll \frac{e(\mathcal{H})^{(r-1)s+1}}{v(\mathcal{H})^{(r-1)s}}$$

of them, $\mathbf{A}_1, \dots, \mathbf{A}_s$ are orderings of the edges of $(r-1)$-stars that form a constellation with base edge A and, as $\{\mathbf{A}_1, \dots, \mathbf{A}_s\}$ and A range over all such pairs and i ranges over $[\![r]\!]$, the sets $E_i(\mathbf{A}_1) \cup \dots \cup E_i(\mathbf{A}_s)$ are distinct edges of $\mathcal{R}_\star^\star$. □

Proof of Lemma 6.6. The lower bound

$$e(\mathcal{M}_\star) \geqslant \left(\frac{2^{-r}}{(r-1)!} - o(1)\right) \cdot \frac{e(\mathcal{H})^{r-1}}{v(\mathcal{H})^{r-2}},$$

is proved analogously to the lower bound on $e(\mathcal{R}_\star)$ in Lemma 6.3. Since an $(r-1)$-star forms a connected hypergraph with $r-1$ edges, we have

$$\Delta_1(\mathcal{M}_\star) \leqslant \Gamma' \cdot \Delta_1(\mathcal{H})^{r-1} \leqslant \Gamma' \cdot \left(K \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}\right)^{r-1} \leqslant \Gamma \cdot \frac{e(\mathcal{M}_\star)}{v(\mathcal{M}_\star)}$$

and, consequently,

$$e(\mathcal{M}_\star) \leqslant \Delta_1(\mathcal{M}_\star) \cdot v(\mathcal{H}) \leqslant \Gamma \cdot \frac{e(\mathcal{H})^{r-1}}{v(\mathcal{H})^{r-2}}$$

where Γ and Γ' are constants that depend only on r and K.

Let $T \subseteq V(\mathcal{H})$ be an arbitrary set of size $t \geqslant 2$. For every sequence $\mathbf{t} = (t_i)_{i=1}^{r-1}$ satisfying $0 \leqslant t_i \leqslant s-1$ for all i, we will bound from above the number of $(r-1)$-stars whose edges intersect T according to $\mathbf{t}$, that is, the ith edge of the star intersects T in t_i vertices (not counting the centre vertex). Since $\sum t_i = t \geqslant 2$ and $s \geqslant 3$, we may assume (by symmetry) that either

(i) $2 \leqslant t_1 \leqslant s-1$ or
(ii) $t_1 = t_i = 1$ for some $i \geqslant 2$.

We may enumerate all stars of the above form as follows:

(1) Choose a labeled partition of T according to the intersection pattern $\mathbf{t}$; there are at most $(r-1)^t$ such partitions.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

(2) Choose the first edge of the star and its centre vertex; there are at most $s \cdot \Delta_{t_1}(\mathcal{H})$ choices.
(3) Choose the remaining $r-2$ edges of the star; since the centre vertex is already fixed, there are at most $\Delta_{t_i+1}(\mathcal{H})$ choices for the ith edge.

This gives

$$\deg_{\mathcal{M}_\star}(T) \leqslant O(1) \cdot \max_{\mathbf{t}} \left\{ \Delta_{t_1}(\mathcal{H}) \cdot \prod_{i=2}^{r-1} \Delta_{t_i+1}(\mathcal{H}) \right\},$$

where the maximum ranges over all sequences $\mathbf{t}$ summing to t and satisfying (i) or (ii) above. Since $\mathcal{H}$ is nonclustered, if (i) holds, then $\Delta_{t_1}(\mathcal{H}) \ll p_{\mathcal{H}}^{t_1-1} \cdot e(\mathcal{H})/v(\mathcal{H})$ and if (ii) holds, then $\Delta_{t_i+1}(\mathcal{H}) \ll p_{\mathcal{H}}^{t_i} \cdot e(\mathcal{H})/v(\mathcal{H})$, as $t_i + 1 = 2 \leqslant s-1$ (we also recall that (13) holds always). Therefore, we may conclude that

$$\deg_{\mathcal{M}_\star}(T) \ll O(1) \cdot \max_{\mathbf{t}} \; p_{\mathcal{H}}^{\sum t_i - 1} \cdot \left(\frac{e(\mathcal{H})}{v(\mathcal{H})} \right)^{r-1} = O\left(p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{M}_\star)}{v(\mathcal{M}_\star)} \right),$$

where the second inequality follows from the lower bound on $e(\mathcal{M}_\star)$ that we established at the beginning. This completes the proof of the lemma. □

Proof of Lemma 6.4. Since a constellation together with its base edge forms a connected hypergraph with $(r-1)s+1$ edges and the number of different rainbow colourings of any given constellation can be bounded by a function of r and s only, there are constants Γ' and Γ that depend only on r, s, and K such that

$$\Delta_1(\mathcal{R}_\star^\star) \leqslant \Gamma' \cdot \Delta_1(\mathcal{H})^{(r-1)s+1} = \Gamma' \cdot \left(K \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \right)^{(r-1)s+1} \leqslant \Gamma \cdot \frac{e(\mathcal{R}_\star^\star)}{v(\mathcal{R}_\star^\star)},$$

where the last inequality follows from Lemma 6.3.

Let $T \subseteq V(\mathcal{H}) \times [\![r]\!]$ be an arbitrary set of size $t \geqslant 2$ and let $T_i = T \cap \big(V(\mathcal{H}) \times \{i\}\big)$. If each of $T_1, \dots, T_r$ is nonempty, then $\deg_{\mathcal{R}_\star^\star}(T) = 0$, so we may assume (by symmetry) that $T_r = \emptyset$. For every sequence

$$\mathbf{t} = \big(t_{i,j} : i \in [\![r-1]\!], j \in [\![s]\!]\big)$$

satisfying $0 \leqslant t_{i,j} \leqslant s-1$ for all i and j, we will bound from above the number of rainbow constellations whose r-rainbow stars intersect the set T according to $\mathbf{t}$, that is, the edge coloured i of the jth star intersects $T_i \subseteq T$ in $t_{i,j}$ vertices. Since $\sum t_{i,j} = t \geqslant 2$ and $s \geqslant 3$, we may assume (by symmetry) that either

(i) $2 \leqslant t_{1,1} \leqslant s-1$ or
(ii) $t_{1,1} = t_{i,j} = 1$ for some $(i,j) \neq (1,1)$.

We may enumerate all constellations of the form described above as follows:

(1) Choose a labeled partition of $T_1, \dots, T_{r-1}$ according to the intersection pattern $\mathbf{t}$; there are at most $[(r-1)s]^t$ such partitions.
(2) Choose the edge coloured 1 of the first star; there are at most $\Delta_{t_{1,1}}(\mathcal{H})$ such edges.
(3) Choose the base edge of the constellation; there are at most $(s - t_{1,1}) \cdot \Delta_1(\mathcal{H})$ choices.
(4) Choose all the $s(r-1)-1$ remaining edges of all the stars forming the constellation one-by-one; since the base edge is already fixed, there are at most $\Delta_{t_{i,j}+1}(\mathcal{H})$ choices for the edge coloured i of the jth star.

This gives

$$\deg_{\mathcal{R}_\star^\star}(T) \leqslant O(1) \cdot \Delta_1(\mathcal{H}) \cdot \max_{\mathbf{t}} \left\{ \Delta_{t_{1,1}}(\mathcal{H}) \cdot \prod_{(i,j) \neq (1,1)} \Delta_{t_{i,j}+1}(\mathcal{H}) \right\},$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

where the maximum ranges over all sequences $\mathbf{t}$ summing to t and satisfying either (i) or (ii) above. Since $\mathcal{H}$ is nonclustered, if (i) holds, then $\Delta_{t_{1,1}}(\mathcal{H}) \ll p_{\mathcal{H}}^{t_{1,1}-1} \cdot e(\mathcal{H})/v(\mathcal{H})$ and if (ii) holds, then $\Delta_{t_{i,j}+1}(\mathcal{H}) \ll p_{\mathcal{H}}^{t_{i,j}} \cdot e(\mathcal{H})/v(\mathcal{H})$, as $t_{i,j}+1 = 2 \leqslant s-1$ (we also recall that (13) holds always). Therefore, we may conclude that

$$\deg_{\mathcal{R}_\star^\star}(T) \ll O(1) \cdot \max_{\mathbf{t}} \; p_{\mathcal{H}}^{\sum t_{i,j}-1} \cdot \left(\frac{e(\mathcal{H})}{v(\mathcal{H})}\right)^{(r-1)s+1} = O\left(p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{R}_\star^\star)}{v(\mathcal{R}_\star^\star)}\right),$$

where the last inequality follows from Lemma 6.3. This completes the proof of the lemma. □

Proof of Lemma 6.5. Suppose that $e\big(\mathcal{R}_\star[C]\big) \geqslant \beta \cdot e(\mathcal{R}_\star)$ and let $W = \pi_1(C)$, where $\pi_1 \colon V(\mathcal{H}) \times [\![r]\!] \to V(\mathcal{H})$ is the projection on the first coordinate. Define a random colouring $\psi \colon W \to [\![r]\!]$ as follows: For every $v \in W$, let $\psi(v)$ be the uniformly random colour $i \in [\![r]\!]$ such that $(v,i) \in C$; in other words, $\psi(v)$ is the uniformly chosen random element of $\pi_2\big(C \cap (\{v\} \times [\![r]\!])\big)$; here, π_2 is the projection of $V(\mathcal{H}) \times [\![r]\!]$ on the second coordinate. Since $\mathcal{R}_\star$ is $(r-1)(s-1)$-uniform, we have

$$\mathbb{E}\big[e\big(\mathcal{R}_\star[\psi]\big)\big] \geqslant r^{-(r-1)(s-1)} \cdot e\big(\mathcal{R}_\star[C]\big) \geqslant \beta r^{-(r-1)(s-1)} \cdot e(\mathcal{R}_\star). \tag{15}$$

From now on, let $\psi \colon W \to [\![r]\!]$ be an arbitrary colouring for which (15) holds without the expectation.

Since Lemma 6.3 supplies a positive constant c such that $e(\mathcal{R}_\star) \geqslant c \cdot e(\mathcal{H})^{r-1}/v(\mathcal{H})^{r-2}$, assumption (A5), see Definition 2.8, assures that

$$e\big(\mathcal{R}_\star^\star[\psi]\big) \geqslant \varepsilon' \cdot \frac{e(\mathcal{H})^{(r-1)s+1}}{v(\mathcal{H})^{(r-1)s}}$$

for some positive constant ε' that depends only on β, r, and the sequence $\mathcal{H}$. Finally, Lemma 6.4 supplies a constant Γ' that depends only on r, s, and K such that $\Delta_1(\mathcal{R}_\star^\star) \leqslant \Gamma' \cdot e(\mathcal{H})^{(r-1)s+1}/v(\mathcal{H})^{(r-1)s+1}$. Consequently,

$$e\big(\mathcal{R}_\star^\star[C]\big) \geqslant e\big(\mathcal{R}_\star^\star[\psi]\big) \geqslant \frac{\varepsilon'}{\Gamma'} \cdot v(\mathcal{H}) \cdot \Delta_1(\mathcal{R}_\star^\star) = \frac{\varepsilon'}{\Gamma' r} \cdot v(\mathcal{R}_\star^\star) \cdot \Delta_1(\mathcal{R}_\star^\star) \geqslant \frac{\varepsilon'}{\Gamma' r} \cdot e(\mathcal{R}_\star^\star),$$

which concludes the proof of the lemma. □

Proof of Lemma 6.7. We have

$$\begin{aligned}\mathrm{Var}'(\mathcal{M}_\star[Z]) &= \sum_{\substack{A,B \in \mathcal{M}_\star \\ A \cap B \neq \emptyset}} p^{|A \cup B|} \leqslant \sum_{A \in \mathcal{M}_\star} p^{|A|} \sum_{\emptyset \neq T \subseteq A} \sum_{B \cap A = T} p^{|B \setminus A|} \\ &\leqslant \mathbb{E}[e(\mathcal{M}_\star[Z])] \cdot \sum_{t=1}^{(r-1)(s-1)} \binom{(r-1)(s-1)}{t} \cdot \Delta_t(\mathcal{M}_\star) \cdot p^{(r-1)(s-1)-t}.\end{aligned}$$

By Lemma 6.6, for some constant Γ' that depends only on r and the sequence $\mathcal{H}$,

$$\begin{aligned}\mathrm{Var}'(\mathcal{M}_\star[Z]) &\leqslant \Gamma' \cdot \mathbb{E}[e(\mathcal{M}_\star[Z])] \cdot \sum_{t=1}^{(r-1)(s-1)} p^{(r-1)(s-1)-t} \cdot p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{M}_\star)}{v(\mathcal{M}_\star)} \\ &= \frac{\Gamma' \cdot \mathbb{E}[e(\mathcal{M}_\star[Z])]^2}{p v(\mathcal{M}_\star)} \cdot \sum_{t=1}^{(r-1)(s-1)} \left(\frac{p_{\mathcal{H}}}{p}\right)^{t-1},\end{aligned}$$

which, by our assumption that $p \geqslant c \cdot p_{\mathcal{H}}$ implies the claimed bound on $\mathrm{Var}'(\mathcal{M}_\star[Z])$. □

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

6.3. *Proof of Theorem 6.1*

We wish to show that, with probability close to one, every colouring of $Z \sim V(\mathcal{H})_p$ that admits only a small number of rainbow constellations will also admit only a small number of rainbow stars. To do this, we will apply the container lemma to the hypergraph $\mathcal{R}^\star_\star$ and conclude that every such colouring is contained in one of $\exp(o(pN))$ subsets of $V(\mathcal{H}) \times [\![r]\!]$, each of which induces $o(e(\mathcal{R}^\star_\star))$ rainbow constellations. By Lemma 6.5, each such container can only induce $o(e(\mathcal{R}_\star))$ rainbow stars. Intuitively, it seems plausible that every colouring of Z residing inside each container should also have a small number of rainbow stars, as we wanted. In order to show this, however, we appear to need an upper bound on the upper tail of the number of rainbow stars in the intersection of $Z \times [\![r]\!]$ with a given container C that is strong enough to survive the union bound over all containers. Unfortunately, the upper tail is most likely too heavy to permit such a naive union bound. We will avert this problem by showing that the overall number of stars in Z is concentrated (for our purposes, a simple second moment argument would do, which is the task of Lemma 6.7) and then transform the question of bounding the upper tail of the number of rainbow stars in $C \cap (Z \times [\![r]\!])$ to that of bounding the lower tail of the number of rainbow stars that are not contained in $C \cap (Z \times [\![r]\!])$; here, Janson's inequality provides an adequate, exponential bound.

Proof of Theorem 6.1. For the sake of brevity, we write V in place of $V(\mathcal{H})$. Let $\Gamma := \Gamma_{6.4}$, let $\beta := \beta_s/(3\Gamma r!)$, let ε be the constant supplied by Lemma 6.5 invoked with $\beta_{6.5} = \beta$, let $c := c_{6.3}$, and let t and δ be the constants from the assertion of the container lemma (Theorem B) invoked with $k_{4.7} = (r-1)(s-1)s$, $K_{4.7} = \Gamma$, and $\varepsilon_{4.7} = \varepsilon$. Further, let $T = T(t, r, \beta, \Gamma)$ be sufficiently large so that

$$\left(\frac{2^t eT}{t}\right)^{tr/T} \leqslant \exp\left(\frac{\beta^2}{5\Gamma}\right) \tag{16}$$

and let $\tau := p/T$. Finally, let $\beta_c := c\delta \cdot T^{-(r-1)(s-1)s}$ and note that

$$\frac{\beta_c}{c} \cdot p^{(r-1)(s-1)s} \leqslant \delta\tau^{(r-1)(s-1)s}. \tag{17}$$

Since Lemma 6.4 implies that, for every $\ell \in [\![(r-1)(s-1)s]\!]$,

$$\Delta_\ell(\mathcal{R}^\star_\star) \leqslant \Gamma \cdot \tau^{\ell-1} \cdot \frac{e(\mathcal{R}^\star_\star)}{v(\mathcal{R}^\star_\star)},$$

Theorem B supplies a function $f\colon \mathcal{P}\big(V \times [\![r]\!]\big)^t \to \mathcal{P}\big(V \times [\![r]\!]\big)$ such that:[9]

(i) For every partial $[\![r]\!]$-colouring ψ with fewer than $\delta\tau^{(r-1)(s-1)s} \cdot e(\mathcal{R}^\star_\star)$ rainbow constellations, there are $S_1, \dots, S_t \subseteq \psi$ with at most $r\tau v(\mathcal{H})$ elements each such that $\psi \subseteq f(S_1, \dots, S_t)$.
(ii) For every $S_1, \dots, S_t \subseteq V \times [\![r]\!]$, the set $f(S_1, \dots, S_t)$ induces fewer than $\varepsilon e(\mathcal{R}^\star_\star)$ edges in $\mathcal{R}^\star_\star$ and thus, by Lemma 6.5, fewer than $\beta e(\mathcal{R}_\star)$ edges in $\mathcal{R}_\star$.

Suppose now that Z fails to satisfy the assertion of the theorem, that is, there is a partial $[\![r]\!]$-colouring ψ of Z such that

$$e\big(\mathcal{R}_\star[\psi]\big) \geqslant \beta_s \cdot \left(\frac{p}{p_{\mathcal{H}}}\right)^{(r-1)(s-1)} \cdot v(\mathcal{H}) = \beta_s \cdot p^{(r-1)(s-1)} \cdot \frac{e(\mathcal{H})^{r-1}}{v(\mathcal{H})^{r-2}}$$

but, nevertheless,

$$e\big(\mathcal{R}^\star_\star[\psi]\big) < \beta_c \cdot \left(\frac{p}{p_{\mathcal{H}}}\right)^{(r-1)(s-1)s} \cdot e(\mathcal{H}) = \beta_c \cdot p^{(r-1)(s-1)s} \cdot \frac{e(\mathcal{H})^{(r-1)s+1}}{v(\mathcal{H})^{(r-1)s}}.$$

[9] Recall that we view partial $[\![r]\!]$-colourings of V as subsets of $V \times [\![r]\!]$.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Consequently, by Lemma 6.6,

$$e(\mathcal{R}_\star[\psi]) \geqslant \frac{\beta_s}{\Gamma} \cdot p^{(r-1)(s-1)} \cdot e(\mathcal{M}_\star) = 3\beta \cdot r! \cdot \mathbb{E}\big[e\big(\mathcal{M}_\star[Z]\big)\big] \tag{18}$$

and, by Lemma 6.3 and (17),

$$e\big(\mathcal{R}^\star_\star[\psi]\big) < \frac{\beta_c}{c} \cdot p^{(r-1)(s-1)s} \cdot e(\mathcal{R}^\star_\star) \leqslant \delta\tau^{(r-1)(s-1)s} \cdot e(\mathcal{R}^\star_\star). \tag{19}$$

Property (i) implies that there are sets $S_1, \dots, S_t \subseteq V \times [\![r]\!]$ with at most $r\tau v(\mathcal{H})$ elements each such that π_1, the projection onto the first coordinate, maps $S_1 \cup \cdots \cup S_t$ injectively to Z and ψ is contained in the set $f(S_1, \dots, S_t)$. In particular, by (18),

$$e\big(\mathcal{R}_\star\big[f(S_1, \dots, S_t) \cap (Z \times [\![r]\!])\big]\big) \geqslant e(\mathcal{R}_\star[\psi]) \geqslant 3\beta \cdot r! \cdot \mathbb{E}\big[e\big(\mathcal{M}_\star[Z]\big)\big]. \tag{20}$$

On the other hand, property (ii) states that $f(S_1, \dots, S_t)$ induces fewer than $\beta e(\mathcal{R}_\star)$ edges in $\mathcal{R}_\star$. We will now show that it is unlikely that this holds for any sequence $S_1, \dots, S_t$.

Let $\mathcal{U}$ be the event that

$$e\big(\mathcal{M}_\star[Z]\big) \leqslant (1+\beta) \cdot \mathbb{E}\big[e\big(\mathcal{M}_\star[Z]\big)\big].$$

It follows from Lemma 6.7 and Markov's inequality that

$$\mathbb{P}(\mathcal{U}^c) \leqslant \frac{4}{\beta^2} \cdot \frac{\mathrm{Var}\big(e(\mathcal{M}_\star[Z])\big)}{\mathbb{E}\big[e(\mathcal{M}_\star[Z])\big]^2} \leqslant \frac{4\Gamma}{\beta^2 \cdot pv(\mathcal{H})} = o(1),$$

as $pv(\mathcal{H}) \to \infty$, see Fact 2.3.

Given a $C \subseteq V \times [\![r]\!]$, define

$$\mathcal{R}^C_\star := \mathcal{R}_\star \setminus \mathcal{R}_\star[C]$$

and let $\mathcal{M}^C_\star$ be the multiset projection of $\mathcal{R}^C_\star$ onto V, so that $e(\mathcal{M}^C_\star) = e(\mathcal{R}^C_\star)$. In particular,

$$\begin{aligned} e\big(\mathcal{R}_\star\big[C \cap (Z \times [\![r]\!])\big]\big) &= e\big(\mathcal{R}_\star\big[Z \times [\![r]\!]\big]\big) - e\Big(\mathcal{R}^C_\star\big[Z \times [\![r]\!]\big]\Big) \\ &= r! \cdot e\big(\mathcal{M}_\star[Z]\big) - e\big(\mathcal{M}^C_\star[Z]\big), \end{aligned} \tag{21}$$

where the final equality holds because every $(r-1)$-star admits exactly $r!$ rainbow colourings (and therefore every edge of $\mathcal{M}_\star[Z]$ corresponds to $r!$ edges of $\mathcal{R}_\star\big[Z \times [\![r]\!]\big]$).

We conclude this discussion with the following observation: If $\mathcal{U}$ holds, then inequality (20) and identity (21) with $C = f(S_1, \dots, S_t)$ imply that

$$e\big(\mathcal{M}^{f(S_1,\dots,S_t)}_\star[Z]\big) \leqslant (1-2\beta) \cdot r! \cdot \mathbb{E}\big[e\big(\mathcal{M}_\star[Z]\big)\big]. \tag{22}$$

Let $\mathfrak{S}$ denote the collection of all sequences $(S_1, \dots, S_t)$ of t subsets of $V \times [\![r]\!]$ with at most $r\tau v(\mathcal{H})$ elements each such that π_1 restricted to $S_1 \cup \cdots \cup S_t$ is injective. Given a sequence $\mathcal{S} = (S_1, \dots, S_t) \in \mathfrak{S}$, let $\mathcal{Y}_\mathcal{S}$ denote the event that $\pi_1(S_1 \cup \cdots \cup S_t) \subseteq Z$ and let $\mathcal{L}_\mathcal{S}$ denote the event that (22) holds; note that $\mathcal{Y}_\mathcal{S}$ is increasing (in Z) whereas $\mathcal{L}_\mathcal{S}$ is decreasing. Recall that, by Harris's inequality [15], increasing and decreasing events in a product probability space are anticorrelated. Therefore, if we let $\mathcal{F}$ denote the event that Z fails the assertion of the theorem, then the above discussion implies that

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

$$\begin{aligned}\mathbb{P}(\mathcal{F}) &\leqslant \mathbb{P}(\mathcal{U}^c) + \sum_{\mathcal{S}\in\mathfrak{S}} \mathbb{P}(\mathcal{L}_{\mathcal{S}} \cap \mathcal{Y}_{\mathcal{S}}) \leqslant \mathbb{P}(\mathcal{U}^c) + \sum_{\mathcal{S}\in\mathfrak{S}} \mathbb{P}(\mathcal{L}_{\mathcal{S}}) \cdot \mathbb{P}(\mathcal{Y}_{\mathcal{S}}) \\ &\leqslant \max_{\mathcal{S}\in\mathfrak{S}} \mathbb{P}(\mathcal{L}_{\mathcal{S}}) \cdot \sum_{\mathcal{S}\in\mathfrak{S}} \mathbb{P}(\mathcal{Y}_{\mathcal{S}}) + o(1).\end{aligned} \tag{23}$$

Claim 6.8. *For every $\mathcal{S} \in \mathfrak{S}$,*

$$\mathbb{P}(\mathcal{L}_{\mathcal{S}}) \leqslant \exp\left(-\frac{\beta^2}{2\Gamma} \cdot pv(\mathcal{H})\right).$$

Proof of Claim 6.8. Fix an arbitrary sequence $\mathcal{S} = (S_1, \dots, S_t) \in \mathfrak{S}$ and recall that

$$\begin{aligned}e\big(\mathcal{M}_\star^{f(S_1,\dots,S_t)}\big) = e\big(\mathcal{R}_\star^{f(S_1,\dots,S_t)}\big) &= e(\mathcal{R}_\star) - e\big(\mathcal{R}_\star\big[f(S_1,\dots,S_t)\big]\big) \\ &\geqslant (1-\beta)\cdot e(\mathcal{R}_\star) = (1-\beta)\cdot r!\cdot e(\mathcal{M}_\star)\end{aligned}$$

or, equivalently,

$$\mathbb{E}\Big[e\big(\mathcal{M}_\star^{f(S_1,\dots,S_t)}[Z]\big)\Big] \geqslant (1-\beta)\cdot r!\cdot \mathbb{E}\big[e\big(\mathcal{M}_\star[Z]\big)\big]. \tag{24}$$

Let $X \coloneqq e\big(\mathcal{M}_\star^{f(S_1,\dots,S_t)}[Z]\big)$ and let $\mu \coloneqq r!\cdot\mathbb{E}\big[e\big(\mathcal{M}_\star[Z]\big)\big]$, so that (22) and (24) can be rewritten as $X \leqslant (1-2\beta)\mu$ and $\mathbb{E}[X] \geqslant (1-\beta)\mu$, respectively. It follows from Janson's inequality (Theorem 4.9) that

$$\mathbb{P}(\mathcal{L}_{\mathcal{S}}) = \mathbb{P}(X \leqslant (1-2\beta)\mu) \leqslant \mathbb{P}(X \leqslant \mathbb{E}[X] - \beta\mu) \leqslant \exp\left(-\frac{\beta^2\mu^2}{2\mathrm{Var}'\big(\mathcal{M}_\star^{f(S_1,\dots,S_t)}[Z]\big)}\right),$$

where, similarly as in Lemma 6.7, we write $\mathrm{Var}'\big(\mathcal{M}_\star^{f(S_1,\dots,S_t)}[Z]\big)$ to denote the pseudo-variance of the sequence of events $(\{A \subseteq Z\})_{A\in\mathcal{M}_\star^{f(S_1,\dots,S_t)}}$. Since $\mathcal{M}_\star^{f(S_1,\dots,S_t)} \subseteq r!\cdot\mathcal{M}_\star$, we have (using an analogous notational convention)

$$\mathrm{Var}'\big(\mathcal{M}_\star^{f(S_1,\dots,S_t)}[Z]\big) \leqslant \mathrm{Var}'(r!\cdot\mathcal{M}_\star[Z]) = (r!)^2\cdot\mathrm{Var}'(\mathcal{M}_\star[Z])$$

Further, by Lemma 6.7,

$$\mathrm{Var}'(\mathcal{M}_\star[Z]) \leqslant \Gamma\cdot\frac{\mathbb{E}[e(\mathcal{M}_\star[Z])]^2}{pv(\mathcal{M}_\star)} = \Gamma\cdot\frac{(\mu/r!)^2}{pv(\mathcal{H})}.$$

Substituting this estimate back into the upper bound on $\mathbb{P}(\mathcal{L}_{\mathcal{S}})$ gives the assertion of the claim. □

Finally, we derive an upper bound on the sum in the right-hand side of (23). To this end, for every integer m, we let

$$\mathfrak{S}_m \coloneqq \big\{(S_1,\dots,S_t)\in\mathfrak{S} : |S_1\cup\cdots\cup S_t| = |\pi_1(S_1\cup\cdots\cup S_t)| = m\big\}$$

and note that

$$\mathfrak{S} = \bigcup_{m=0}^{tr\tau v(\mathcal{H})} \mathfrak{S}_m.$$

It is not hard to see that, for every m,

$$|\mathfrak{S}_m| \leqslant \binom{v(\mathcal{H})}{m}\cdot\big(2^t r\big)^m \leqslant \left(\frac{2^t e r v(\mathcal{H})}{m}\right)^m$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

and $\mathbb{P}(\mathcal{Y}_{\mathcal{S}}) = p^m$ for every $\mathcal{S} \in \mathfrak{S}_m$. Consequently, since $\tau \leqslant p$ and, for every positive real a, the function $x \mapsto (ea/x)^x$ is increasing on the interval $[0, a]$,

$$\sum_{\mathcal{S}\in\mathfrak{S}} \mathbb{P}(\mathcal{Y}_{\mathcal{S}}) = \sum_{m=0}^{tr\tau v(\mathcal{H})} \left(\frac{2^t erpv(\mathcal{H})}{m}\right)^m \leqslant v(\mathcal{H}) \cdot \left(\frac{2^t ep}{t\tau}\right)^{tr\tau v(\mathcal{H})}$$
$$= v(\mathcal{H}) \cdot \left(\frac{2^t eT}{t}\right)^{(tr/T)\cdot pv(\mathcal{H})} \overset{(16)}{\leqslant} \exp\left(\frac{\beta^2}{4\Gamma} \cdot pv(\mathcal{H})\right).$$

We may finally conclude that

$$\mathbb{P}(\mathcal{F}) \leqslant o(1) + \exp\left(\frac{\beta^2}{4\Gamma} \cdot pv(\mathcal{H}) - \frac{\beta^2}{2\Gamma} \cdot pv(\mathcal{H})\right) = o(1),$$

where we again used the assumption that $pv(\mathcal{H}) \to \infty$, see Fact 2.3. □

7. Applications

In this section, we will use our general Theorem 2.1 to prove Theorems 1.2, 1.3, and 1.4. In the following three subsections, we will verify that Theorem 2.1 can be applied to hypergraphs that naturally arise in the context of Ramsey questions for: graphs (Section 7.1), arithmetic progressions (Section 7.2), and Schur triples (Section 7.3). In each case, we will verify the list of assumptions of the theorem, which will swiftly award us with a sharp threshold result for the corresponding Ramsey problem. As a reminder, the assumptions on the hypergraph $\mathcal{H}$ are:

(A1) is symmetric,
(A2) it is nonclustered,
(A3) noncolourability has a threshold at $p_{\mathcal{H}}$,
(A4) it satisfies 2-choosability for typical bounded-sized subsets, and
(A5) it satisfies the rainbow star-constellation property.

The first two assumptions are more technical in nature and will mostly set the framework of the application. For example, in the arithmetic setting, (A1) will force us to work in $\mathbb{Z}_N$ instead of $[\![N]\!]$ whereas, for graphs, (A2) is akin to requiring the graph to be strictly 2-balanced. Assumption (A3) was the subject of previous works. Having said that, in the arithmetic setting of Theorems 1.3 and 1.4, the 0-statements implicit in (A3) are marginally stronger than what was explicitly established by previous works, due to the fact that we are working in modular arithmetic. Even though these earlier works can be adapted to yield (A3), we will not dwell on it and instead establish stronger forms of these 0-statements that also imply (A4), Theorems 1.5 and 1.6. These follow from the general Theorem 7.4, which will be proved in Section 7.4. (We also recall that the 1-statements can be derived using Proposition 2.11 since the hypergraphs in all of our applications are robustly noncolourable.) As for (A5), even though it is rather straightforward to verify in the context of Theorem 1.3, checking it in the remaining two cases is far from easy. (We recall here that this assumption is not satisfied in all cases of interest, see Appendix C.) Summarising, the bulk of the work in this section will be spent in verifying the last two assumptions.

7.1. Graphs

In this section, we will prove Theorem 1.2, which asserts that the property $G_{n,p} \to (H)_r$ has a sharp threshold for certain pairs of H and r. Given a strictly 2-balanced graph H, let $\mathcal{H}_H$ be the hypergraph of copies of H in K_n whose vertices are the edges of K_n and whose hyperedges are (the edge sets of) all copies of H. It is straightforward to check that $s = e_H$, $v(\mathcal{H}_H) = \Theta(n^2)$, $e(\mathcal{H}_H) = \Theta(n^{v_H})$, and $p_{\mathcal{H}_H} = \Theta(n^{-1/m_2(H)})$. Moreover, it is routine to verify that $\mathcal{H}_H$ is nonclustered (as H is strictly 2-balanced) and that it is symmetric (due to the symmetries of K_n). It was proved by Rödl and Ruciński

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

[25] that, for each $r \geqslant 2$, the threshold for non-r-colourability of $(\mathcal{H}_H)_p$ is located at $p_{\mathcal{H}}$. Thus, assumptions (A1)– (A3) are met. We now turn to verifying the last two assumptions.

7.1.1. Rainbow star-constellation property for graphs

A rainbow star in the hypergraph $\mathcal{H}_H$ of copies of a graph H in K_n is comprised of $r-1$ monochromatic copies of H minus some edge (not necessarily the same edge in different copies), each coloured with a different colour, that are glued on that missing edge, called the centre of the star. A rainbow constellation is a union of edge-disjoint rainbow stars whose centres form a copy of H. We will tacitly assume that all stars and constellations are generic in the sense that the copies of H minus an edge that form them do not share more vertices than necessary: every rainbow star will have $v_S := (v_H - 2) \cdot (r-1) + 2$ vertices and every rainbow constellation will have $v_C := e_H \cdot (v_H - 2) \cdot (r-1) + v_H$ vertices.

To prove the rainbow star-constellation property in r colours we will need to show that any collection of $\Theta(e(\mathcal{H}_H)^{r-1}/v(\mathcal{H}_H)^{r-2}) = \Theta(n^{v_S})$ many $(r-1)$-stars, induces $\Theta(e(\mathcal{H}_H)^{s(r-1)+1}/v(\mathcal{H}_H)^{s(r-1)}) = \Theta(n^{v_C})$ many $(r-1)$-constellations. Put differently, if we have a collection achieving the full count of rainbow stars in K_n, then it induces the full count of rainbow constellations.

Let F and G be edge-coloured graphs. A homomorphism from F to G is a function $\varphi\colon V(F) \to V(G)$ such that, for every $uv \in F$, the pair $\{\varphi(u), \varphi(v)\}$ is an edge of G with the same colour as uv. We wish to prove the following characterisation for the rainbow star-constellation property.

Proposition 7.1. *The hypergraph $\mathcal{H}_H$ of copies of a given graph H in K_n has the rainbow star-constellation property if and only if every rainbow star S of H admits a rainbow constellation C of H with a homomorphism $C \to S$.*

In order to prove this statement, we will first gather some definitions and tools. Given edge-coloured graphs F and G, let $\mathrm{Hom}(F,G)$ denote the family of all homomorphisms from F to G and let $\hom(F,G)$ denote the density of $\mathrm{Hom}(F,G)$ in $V(G)^{V(F)}$, that is,

$$\hom(F,G) := \frac{|\mathrm{Hom}(F,G)|}{v_G^{v_F}}.$$

In other words, $\hom(F,G)$ is the probability that a random function $\varphi\colon V(F) \to V(G)$ belongs to $\mathrm{Hom}(F,G)$.

Observation 7.2. Let F and G be graphs. If $F' \subseteq F$, then $\hom(F',G) \geqslant \hom(F,G)$.

Proof. Indeed, if $\varphi \in \mathrm{Hom}(F,G)$, then $\varphi|_{V(F')} \in \mathrm{Hom}(F',G)$. Therefore, $|\mathrm{Hom}(F,G)| \leqslant |\mathrm{Hom}(F',G)| \cdot v_G{}^{v_F - v_{F'}}$ and the result follows. □

For a graph F and a positive integer k, the k-blowup of F, denoted by $F^{(k)}$, is the graph obtained from F by replacing each vertex a of F by an independent set V_a of size k and every edge ab of F by a complete bipartite graph between V_a and V_b.

Observation 7.3. $\mathrm{Hom}(F,G) \neq \emptyset$ if and only if $F \subseteq G^{(v_F)}$.

The following lemma should be folklore.

Lemma 7.4. *Suppose that F and G are graphs and let k be a positive integer. Then*

$$\hom(F^{(k)},G) \geqslant \hom(F,G)^{k \cdot v_F}.$$

Proof. Let $v = v_F$ and let $u_1, \ldots, u_v$ be an arbitrary ordering of the vertices of F. For $i \in \{0, \ldots, v\}$ let F_i denote the graph obtained from F by blowing up vertices $u_1, \ldots, u_i$ by a factor of k, so that $F_0 = F$ and $F_v = F^{(k)}$. It suffices to show that, for every $i \in [\![v]\!]$,

$$\hom(F_i,G) \geqslant \hom(F_{i-1},G)^k.$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Let $\varphi\colon V(F_i) \to V(G)$ be a random function, let $u_{i,1}, \dots, u_{i,k}$ be the k copies of u_i in F_i, and let $V_i' = V(F_i) \setminus \{u_{i,1}, \dots, u_{i,k}\} = V(F_{i-1}) \setminus \{u_i\}$. We have

$$\begin{aligned}\hom(F_i, G) = \mathbb{P}\big(\varphi \in \mathrm{Hom}(F_i, G)\big) &= \mathbb{P}\big(\varphi|_{V_i'} \in \mathrm{Hom}(F_i[V_i'], G)\big) \\ &\quad \cdot \mathbb{P}\big(\varphi \in \mathrm{Hom}(F_i, G) \mid \varphi|_{V_i'} \in \mathrm{Hom}(F_i[V_i'], G)\big)\end{aligned}$$

Since $F_i[V_i'] = F_{i-1}[V_i']$, we have

$$\mathbb{P}\big(\varphi|_{V_i'} \in \mathrm{Hom}(F_i[V_i'], G)\big) = \mathbb{P}\big(\varphi|_{V_i'} \in \mathrm{Hom}(F_{i-1}[V_i'], G])\big) = \hom(F_{i-1}[V_i'], G).$$

Crucially, since $\varphi(u_{i,1}), \dots, \varphi(u_{i,k})$ are independent, uniformly random elements of $V(G)$ and $\{u_{i,1}, \dots, u_{i,k}\}$ is an independent set in F_i,

$$\begin{aligned}&\mathbb{P}\big(\varphi \in \mathrm{Hom}(F_i, G) \mid \varphi|_{V_i'} \in \mathrm{Hom}(F_i[V_i'], G)\big) \\ &\qquad = \prod_{j=1}^{k} \mathbb{P}\big(\varphi|_{V_i' \cup \{u_{i,j}\}} \in \mathrm{Hom}(F_i[V_i' \cup \{u_{i,j}\}], G) \mid \varphi|_{V_i'} \in \mathrm{Hom}(F_i[V_i'], G)\big).\end{aligned}$$

Finally, since $F_i[V_i' \cup \{u_{i,j}\}] \cong F_{i-1}$, we conclude that, letting $\psi\colon V(F_{i-1}) \to V(G)$ be a uniformly random map,

$$\begin{aligned}\hom(F_i, G) &= \hom(F_{i-1}[V_i'], G) \cdot \mathbb{P}\big(\psi \in \mathrm{Hom}(F_{i-1}, G) \mid \psi|_{V_i'} \in \mathrm{Hom}(F_{i-1}[V_i'], G)\big)^k \\ &\geqslant \hom(F_{i-1}[V_i'], G)^k \cdot \mathbb{P}\big(\psi \in \mathrm{Hom}(F_{i-1}, G) \mid \psi|_{V_i'} \in \mathrm{Hom}(F_{i-1}[V_i'], G)\big)^k \\ &= \hom(F_{i-1}, G)^k,\end{aligned}$$

as claimed. □

Proof of Proposition 7.1. First, suppose that, for some rainbow star S, there is no constellation C such that $C \to S$. The blowup $S^{(\lfloor n/v_S \rfloor)}$ is an edge-coloured subgraph of K_n with $\Omega(n^{v_S})$ rainbow stars but no rainbow constellations, see Observation 7.3.

For the other direction, suppose that G is an edge-coloured subgraph of K_n that has $\Omega(n^{v_S})$ rainbow stars. Since there are only $O(1)$ isomorphism types of stars of H, there must be some rainbow star S of H such that $\hom(S, G)$ is bounded from below by a positive constant. Let C be a rainbow constellation satisfying $C \to S$, that is, $C \subseteq S^{(v_C)}$, see Observation 7.3. By Observation 7.2 and Lemma 7.4,

$$\hom(C, G) \geqslant \hom(S^{(v_C)}, G) \geqslant \hom(S, G)^{v_C \cdot v_S} = \Omega(1).$$

In particular, G contains $\Omega(n^{v_C})$ rainbow constellations. □

We will now derive several sufficient conditions for the rainbow star-constellation property that are easier to verify than the abstract criterion provided by Proposition 7.1.

Corollary 7.5. *If H is bipartite, then $\mathcal{H}_H$ has the rainbow star-constellation property.*

Proof. We wish to show that every rainbow star S of a bipartite graph admits a rainbow constellation C such that $C \to S$. Given a star S, we will build the required constellation C in the following way. Begin with a copy of H with partite sets U and V. For all $u \in U$ and $v \in V$ that are adjacent in H, place a copy of S centred at uv. It is not hard to check that the function that maps every vertex in each copy of S in C to its corresponding vertex of S is a homomorphism from C to S. □

Next, observe that the rainbow star-constellation property for r colours is monotone decreasing in r. However, as it turns out, whenever it holds with $r = 3$, it also holds for all r strictly greater than three. We will show this by describing another equivalent property that does not mention rainbow structures and instead deals solely with the symmetries of H.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Definition 7.6. A graph H is *collapsible* if, for every edge $e \in H$ and every vertex $a \in e$, there is an edge $f \in H$ and a homomorphism $H \setminus f \to H \setminus e$ mapping both endpoints of f to a.

Corollary 7.7. *Suppose that $r > 2$ and that H is not bipartite. Then $\mathcal{H}_H$ has the rainbow star-constellation property if and only if H is collapsible.*

Proof. Suppose first that $\mathcal{H}_H$ has the rainbow star-constellation property for $r > 2$. Let $e = \{a, b\}$ be an edge of H. We will construct a rainbow star S centred at some uv using copies of $H \setminus e$ as follows: In the first colour, use a copy of $H \setminus e$ so that u plays the role of a and v plays the role of b. In every other colour (the number of colours is $r - 1 \geqslant 2$), we switch the roles – u plays the role of b and v plays the role of a. By Proposition 7.1, we are guaranteed a rainbow constellation C and a homomorphism φ from C to S. Let W be the set of vertices spanned by the centres of stars comprising C. We claim that $\varphi(W) \subseteq \{u, v\}$. Indeed, every vertex in W is incident to edges of all $r - 1 > 1$ colours and every vertex of S other then u and v is incident to edges of solely one colour. Since W is the vertex set of a copy of H whose every edge is the centre of a star in C, the assumption that H is not bipartite implies that φ maps both endpoints of at least one such centre to the same vertex (either u or v). Let S' be the star in C whose centre has this property. The restriction of φ to S' (specifically, the first two colours of S') describes two homomorphisms from copies of H minus an edge, one of them mapping both endpoints of the missing edge to a, and the other mapping both of them to b.

Suppose now that H is collapsible. Given a rainbow star S, centred at some uv, there are edges $a_1b_1, \ldots, a_{r-1}b_{r-1} \in H$ such that S comprises copies of $H \setminus a_1b_1, \ldots, H \setminus a_{r-1}b_{r-1}$ glued together so that $a_1, \ldots, a_{r-1}$ are all mapped to u and $b_1, \ldots, b_{r-1}$ are all mapped to v. For each $i \in [\![r-1]\!]$, we may find an $f_i \in H$ and a homomorphism φ_i from $H \setminus f_i$ to $H \setminus a_ib_i$ that maps both endpoints of f_i to a_i. Let S' be a rainbow star comprised of copies of $H \setminus f_1, \ldots, H \setminus f_{r-1}$ glued together along $f_1, \ldots, f_{r-1}$. It is not hard to check that the function φ' extending all $\varphi_1, \ldots, \varphi_{r-1}$ is a homomorphism from S' to S that maps both endpoints of the centre of S' to u. Finally, let C' be a rainbow constellation whose each star is a copy of S'. The function mapping each vertex in each copy of S' according to φ' is a homomorphism from C' to S (which maps all vertices spanned by the base of C' to u). □

Last, in the case where $r = 2$ the notion of rainbow stars and constellations degenerates – a rainbow star is just a copy of H minus an edge. As a result, we have fewer restrictions than in the case $r > 2$ and even a weaker version of the collapsibility will be sufficient to imply the star-constellation property.

Definition 7.8. A graph H is *semi-collapsible* if, for every edge $e \in H$, there is an edge $f \in H$ and a homomorphism $H \setminus f \to H \setminus e$ mapping both endpoints of f to the same vertex.

Corollary 7.9. *If $r = 2$ and H is semi-collapsible, then $\mathcal{H}_H$ has the star-constellation property.*

Proof. Suppose that H is semi-collapsible and let $H \setminus e$, where $e \in H$, be an arbitrary star of H. We let C be a constellation constructed from copies of $H \setminus f$, for some $f \in H$ such that there exists a homomorphism φ from $H \setminus f$ to $H \setminus e$ that maps both endpoints of f to the same vertex. It is not hard to check that the function that maps each vertex in each copy of $H \setminus f$ in C according to φ is a homomorphism from C to $H \setminus e$. □

One can use these corollaries to establish the rainbow star-constellation property for certain families of graphs. We will give two examples. Recall that a graph is called nearly bipartite if it can be made bipartite by removing one edge; for example, a cycle. Together with trees, strictly 2-balanced nearly bipartite graphs make the largest family of graphs for which sharpness was previously established (albeit, only when $r = 2$). We will show that nearly bipartite graphs are collapsible and so their respective hypergraphs $\mathcal{H}_H$ all have the rainbow star-constellation property for every r. Secondly, we will note that cliques are also collapsible, thus giving another application to a natural family on the other end of the spectrum. We should also note that some graphs do not have the rainbow star-constellation property. One such example is the Petersen graph (see Appendix C).

Corollary 7.10. *If H is nearly bipartite and with minimum degree at least two, then $\mathcal{H}_H$ has the rainbow star-constellation property.*

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Proof. By Corollary 7.5, we may assume that H is not bipartite. By Corollary 7.7, it suffices to show that H is collapsible. Let e be an edge of H and let u be an endpoint of e. Since $\delta(H) \geqslant 2$, there is an edge $e' \neq e$ that also contains u. Now take an edge f such that $H \setminus f$ is bipartite. Consider some bipartition of $H \setminus f$ and let U be the colour class that contains both endpoints of f (we assumed that H is not bipartite). We map $H \setminus f$ to $H \setminus e$ by sending all the vertices in U to u and by sending the vertices in $V(H) \setminus U$ to the second endpoint of e'. □

Corollary 7.11. *If H is a clique, then $\mathcal{H}_H$ has the rainbow star-constellation property.*

Proof. This follows from Corollary 7.7 and the fact that a clique is collapsible. Indeed, given any edge e of the clique and an endpoint u of e, the mapping of $K_n \setminus e$ to itself that maps both endpoints of e to u and fixes every other vertex is a homomorphism. □

7.1.2. List colouring graphs

We wish to show that, when $p \leqslant n^{-1/m_2(H)}$, a.a.s. every constant-sized subgraph of $G_{n,p}$ is *2-choosable with respect to H*, that is, its edges can be coloured from arbitrary lists of size two without introducing a monochromatic copy of H. Since a.a.s. any constant-sized subgraph $G \subseteq G_{n,p}$ satisfies $m(G) \leqslant m_2(H)$, it suffices to prove the following proposition.

Proposition 7.12. *Let H be a graph that is not a forest. If a graph G satisfies $m(G) \leqslant m_2(H)$, then it is 2-choosable with respect to H.*

Before we prove this, let us state a helpful lemma. Given a graph H and a set $W \subseteq V(H)$, it will be convenient to denote by $\bar{e}_H(W)$ the number of edges incident with a vertex of W, that is, $\bar{e}_H(W) \coloneqq e_H - e_{H\setminus W}$.

Lemma 7.13 (Helpful Lemma). *Suppose that H is strictly 2-balanced and suppose that $W \subseteq V(H)$ satisfies $1 \leqslant |W| \leqslant v_H - 3$. Then*

$$m_2(H) \cdot |W| < \bar{e}_H(W).$$

Proof. Since $H - W$ is a proper subgraph of H with at least three vertices and H is strictly 2-balanced,

$$\frac{e_H - 1 - \bar{e}_H(W)}{v_H - 2 - |W|} = \frac{e_{H-W} - 1}{v_{H-W} - 2} < m_2(H) = \frac{e_H - 1}{v_H - 2},$$

which means that $(e_H - 1) \cdot |W| < (v_H - 2) \cdot \bar{e}_H(W)$ and the result follows. □

We now turn to proving the proposition. We will have separate arguments for three small graphs – specifically K_3, K_4, and C_4 – and then a general argument for every other graph. We will begin with the general argument and subsequently supply the remaining cases.

Proof of Proposition 7.12 (Part I: The general argument). Without loss of generality, we may assume that H is strictly 2-balanced. Indeed, otherwise we replace H with one of its minimal subgraphs H' satisfying $m_2(H') = m_2(H) > 1$. Since the only strictly 2-balanced graphs with fewer than five vertices (and a cycle) are K_3, K_4, and C_4, which will require a separate argument, we may further assume that H has at least five vertices.

Suppose that the statement is false and let G be a minimal counterexample. Write $m_2(H) = k + \varepsilon$, where k is an integer and $\varepsilon \in [0, 1)$. Since H is not a forest, we must have $k \geqslant 1$.

The idea of the proof is to use the upper bound on the density of G to locate a sparse subgraph $S \subseteq G$; this will be achieved using a discharging argument. By the minimality of G, we are able to colour $G \setminus S$ without monochromatic copies of H and because S is sufficiently sparse, we will be able to extend each such colouring to all of G.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Claim 7.14. *One of the following holds:*

(1) *G has a vertex of degree at most $2k$,*
(2) *G has a vertex of degree $2k+1$ with a neighbour of degree at most $2k+2$ and $\varepsilon \geqslant 1/2$, or*
(3) *G has a vertex of degree $2k+3$ with two neighbours of degree $2k+1$ and $\varepsilon \geqslant 7/8$.*

Proof. Since $\delta(G) \leqslant 2m(G) \leqslant 2m_2(H) = 2k + 2\varepsilon < 2k+2$, we have $\delta(G) \leqslant 2k+1$ and, if $\varepsilon < 1/2$, then $\delta(G) \leqslant 2k$. We may thus assume that $\delta(G) = 2k+1$ and $\varepsilon \geqslant 1/2$, since otherwise (1) holds. We may further assume that all neighbours of every vertex of degree $2k+1$ have degrees at least $2k+3$, since otherwise (2) holds.

Assign to each $v \in V(G)$ a charge of $\deg(v) - 2(k+\varepsilon)$. Note that the average charge is at most $2m(G) - 2(k+\varepsilon) \leqslant 0$. We define the following discharging rule: every vertex of degree $2k+1$ takes a charge of $\frac{2\varepsilon-1}{2k+1}$ from each of its neighbours. By our assumption, no vertex of degree $2k+1$ or $2k+2$ sends charge to any of its neighbours. In particular, the final charge of a vertex of degree $2k+1$ is

$$2k+1-2(k+\varepsilon)+(2k+1)\cdot\frac{2\varepsilon-1}{2k+1}=0$$

and the final charge of a vertex of degree $2k+2$ is $2k+2-2(k+\varepsilon) > 0$.

Since the total charge remains unchanged, the final charge of some vertex of degree at least $2k+3$ must be nonpositive. Let v be one such vertex. Suppose that $\deg(v) = 2k+t$, where $t \geqslant 3$, and that v has x neighbours with degree $2k+1$. Since the final charge of v is

$$2k+t-2(k+\varepsilon)-x\cdot\frac{2\varepsilon-1}{2k+1}\leqslant 0,$$

we have

$$(t-2)(2k+1) < \frac{t-2\varepsilon}{2\varepsilon-1}\cdot(2k+1)\leqslant x\leqslant 2k+t,$$

which implies that $t < 3 + 1/k \leqslant 4$. Therefore, $t = 3$ and $x > 2k+1 \geqslant 3$, which means that some vertex of degree $2k+3$ has more than three neighbours of degree $2k+1$. Moreover, we also have $\frac{3-2\varepsilon}{2\varepsilon-1}\cdot(2k+1)\leqslant 2k+3$, which implies that $\varepsilon \geqslant \frac{4k+3}{4k+4} \geqslant \frac{7}{8}$. □

We split the argument into three cases, depending on which item in Claim 7.14 holds.

Case 1. Item (1) in Claim 7.14 holds. Since the Helpful Lemma implies that $\deg(a) = \bar{e}_H(\{a\}) > m_2(H) = k+\varepsilon$ for every vertex $a \in V(H)$, we have $\delta(H) \geqslant k+1$. Let v be a vertex of smallest degree in G and let S comprise all the edges incident with v. Any colouring of $G \setminus S$ may be extended to G in the following way: Because all lists have two colours, we may choose colours for the edges of S so that every colour is selected at most $\lceil |S|/2 \rceil = \lceil \delta(G)/2 \rceil = k$ times. Since $\delta(H) > k$, this means that v cannot belong to a monochromatic copy of H.

Case 2. Item (2) in Claim 7.14 holds. Let uv be an arbitrary edge of G satisfying $\deg(u) = 2k+1$ and $\deg(v) \leqslant 2k+2$. We claim that any colouring of $G \setminus uv$ can be extended to G. Suppose that it cannot. This means that uv completes a monochromatic copy of H in both its colour options. But this means that, for some $ab \in H$,

$$4k+1 \geqslant \deg(u)+\deg(v)-2 = \bar{e}_G(\{u,v\})-1 \geqslant 2\cdot\big(\bar{e}_H(\{a,b\})-1\big).$$

However, the Helpful Lemma implies that $\bar{e}_H(\{a,b\}) > 2m_2(H) = 2k+2\varepsilon \geqslant 2k+1$, a contradiction.

Case 3. Item (3) in Claim 7.14 holds. Let u, v_1, and v_2 be distinct vertices of G satisfying $uv_1, uv_2 \in G$, $\deg(u) = 2k+3$, and $\deg(v_1) = \deg(v_2) = 2k+1$. We may assume that $v_1v_2 \notin G$, since otherwise (2) in Claim 7.14 holds. We claim that any proper colouring of $G' := G \setminus \{uv_1, uv_2\}$ can be extended to G. To this end, let $G_1 := G' \cup \{uv_1\}$ and $G_2 := G' \cup \{uv_2\}$. Since $\deg_{G_i}(u) = 2k+2$ and $\deg_{G_i}(v_i) = 2k+1$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

for both $i \in \{1,2\}$, the argument presented in Case 2 shows that any colouring of G' can be (separately) extended to both G_1 and G_2. If some extensions of the colouring of G' to G_1 and G_2 assign different colours to uv_1 and uv_2, then their common extension is an H-free colouring of G. Therefore, we may assume that there is a colour j such that, for both $i \in \{1,2\}$, every extension of the colouring of G' to G_i assigns this colour j to uv_i. This means that uv_i completes a copy of H in G_i whose all remaining edges are assigned a colour other than j (the second colour from the list of uv_i) in the colouring of G'. In particular, there are vertices a_1, a_2, and b_2 of H such that $a_2b_2 \in H$ and, letting m denote the number of edges of G' incident to $\{u, v_1, v_2\}$ that are not coloured j,

$$m \geqslant \bar{e}_H(\{a_1\}) - 1 + \bar{e}_H(\{a_2, b_2\}) - 1$$

As before, the Helpful Lemma implies that $\bar{e}_H(\{a_1\}) > k$ and $\bar{e}_H(\{a_2, b_2\}) > 2k + 2\varepsilon > 2k + 1$, which means that $m \geqslant 3k + 1$. It follows that the number m' of edges of G' incident to $\{u, v_1, v_2\}$ that are coloured j satisfies

$$m' = \bar{e}_{G'}(\{u, v_1, v_2\}) - m = \deg(u) + \deg(v_1) + \deg(v_2) - 4 - m \leqslant 3k.$$

Assign the colour j to both uv_1 and uv_2. If this is not a proper colouring of G, then there must be a copy of H in colour j that contains both uv_1 and uv_2. This means that there are vertices a, b, and c of H such that $ab, bc \in H$, $ac \notin H$, and

$$\bar{e}_H(\{a, b, c\}) - 2 \leqslant m' \leqslant 3k.$$

However, since $\varepsilon \geqslant 1-1/8$, we must have $v_H \geqslant 8+2$ and the Helpful Lemma implies that $\bar{e}_H(\{a, b, c\}) > 3m_2(H) = 3k + 3\varepsilon > 3k + 2$, a contradiction. □

Proof of Proposition 7.12 (Part II: the small graphs). All that remains is to prove the statement for C_4, K_4, and K_3. Assume that G is a minimal counterexample.

(C_4) Since $m(G) \leqslant m_2(C_4) = 3/2$, then either $\delta(G) < 3$ or G is 3-regular. However, since $\delta(C_4) = 2$, the former leads to a contradiction: we would take a vertex v of G of degree at most two, find a C_4-free colouring for $G \setminus \{v\}$, and extend this colouring to G by choosing different colours for the two edges incident with v. The same argument works in the case where $\deg(v) = 3$ and the lists of colours for the three edges incident to v are not identical (we can still choose a different colour for each of these edges). Thus, we may assume that G is 3-regular and all the colour lists are identical (the minimality of G implies that it is connected) – they all contain the colours red and blue.

Let v be an arbitrary vertex of G and let u_1, u_2, and u_3 be its neighbours. Fix a colouring of $G - v$. Since the colouring cannot be extended to G, every pair of vertices among $\{u_1, u_2, u_3\}$ is connected in $G - v$ by both a red and a blue path of length two. (Indeed, if u_1 and u_2 were not connected by a red path, say, then colouring vu_1 and vu_2 red and vu_3 blue would yield a C_4-free colouring.) However, since G is 3-regular, each u_i is incident to at most one red and at most one blue edge of G'. This means that there are vertices $v_r, v_b \neq v$ such that v_r is connected to all u_i in red and v_b is connected to all u_i in blue. Since G is 3-regular, $\{v, v_r, v_b, u_1, u_2, u_3\}$ is a connected component of G; by minimality, $G = G[\{v, v_r, v_b, u_1, u_2, u_2\}] \cong K_{3,3}$. However, $K_{3,3}$ has many 2-edge-colourings without a monochromatic C_4 (e.g., $K_{3,3}$ can be decomposed into C_6 and $3K_2$).

(K_4) Since $m(G) \leqslant m_2(K_4) = 5/2$, then either $\delta(G) \leqslant 4$ or G is 5-regular. Since $\delta(K_4) = 3$, if G has a vertex v of degree at most four or a vertex of degree five whose incident edges have nonidentical colour lists, we may extend any K_4-free colouring of $G - v$ to G by colouring edges incident to v in such a way that every colour is used at most twice. Thus, we may assume that G is 5-regular and all the colour lists contain the colours red and blue.

Let v be an arbitrary vertex and fix a colouring of $G - v$. Since $G - v$ has no red K_4's, there must be a 3-element subset $T \subseteq N(v)$ that does not induce a red triangle. Colouring

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

the three edges connecting v to T red and the remaining two edges incident to v blue yields a K_4-free colouring of G.

(K_3) Since $m(G) \leqslant m_2(K_3) = 2$, then either $\delta(G) \leqslant 3$ or G is 4-regular. Suppose that, for some $v \in V(G)$, there was an orientation of the edges of $G[N(v)]$ in which every vertex had out-degree at most one. We could then extend every K_3-free colouring of $G - v$ to G as follows: For every $u \in N(v)$, the edge uv gets a colour that is different from the colour of the out-edge from u. (Since each triangle involving v contains an edge of $G[N(v)]$, this colouring is K_3-free.) As every graph with at most four edges has such an orientation, we may assume that $e(N(v)) \geqslant 5$ for every $v \in V(G)$; in particular, $\delta(G) \geqslant 4$, so G is 4-regular.

We claim that $G = K_5$. If $e(N(v)) > 5$ for some $v \in V(G)$, then $G = K_5$, since G is 4-regular and connected. We may thus further assume that $e(N(v)) = 5$ for every v. Pick some v and denote $N(v) = \{u_1, u_2, u_3, u_4\}$ so that $u_1u_3 \notin G[N(v)]$. Since G is 4-regular, there must be a $w \in V(G) \setminus (\{v\} \cup N(v))$ such that $N(u_1) = \{v, w, u_2, u_4\}$. Moreover, w is not adjacent to either of v, u_2, and u_4, as they all have 4 neighbours in $\{v\} \cup N(v)$, and thus $e(N(u_1)) \leqslant 3$, a contradiction.

Finally, we show how that K_5 is 2-choosable with respect to K_3. If some colour, say red, contains a 5-cycle, then we may colour this 5-cycle red and the complementary 5-cycle not red. If some colour class, say red, contains an edge, say e, not in a triangle, then we may colour $K_5 \setminus e$ without monochromatic triangles (this is possible as K_5 is minimally non-2-choosable) and colour e red. If none of the above is true, then each colour induces one of the following graphs: K_3, K_4, K_4^-, $K_5 \setminus K_3$, or two triangles sharing a vertex. If some colour, say red, induces $K_5 \setminus K_3$, then we colour $K_{2,3}$ with red, the remaining edge of $K_5 \setminus K_3$ with not red and the edges of the K_3 in the complement with two different colours other than red. If one of the colours, say red, induces K_4 or K_4^-, then colour a C_4 with red and its diagonal with a colour other than red. Each of the remaining, uncoloured four edges can close at most one monochromatic triangle, as red is not available anywhere outside of the K_4 we have already coloured; thus we may colour them one-by-one. This leaves the case where every colour class is either K_3 or two triangles sharing a vertex. But this is impossible, since 3 does not divide $2e(G) = 20$. □

7.2. *Arithmetic progressions*

In this section, we prove Theorem 1.3, which asserts that $\mathcal{W}(k, r)$ has a sharp threshold in $(\mathbb{Z}_N)_p$. Even though the most natural setting for van der Waerden's theorem is the interval $[\![N]\!]$, the corresponding hypergraph lacks the required symmetry. As a result, we have to work in $\mathbb{Z}_N$ instead, where transitivity is guaranteed by translations. (We note that this was also the case in [10], which proved sharpness of the threshold for van der Waerden's theorem in two colours.) We will consider the k-uniform hypergraph $\mathcal{H}_{k\text{-AP}}$ of proper k-term arithmetic progressions in $\mathbb{Z}_N$. It is easily verified that $\mathcal{H}_{k\text{-AP}}$ has $\Theta(N^2)$ edges, and thus $p_{\mathcal{H}_{k\text{-AP}}} = \Theta(N^{-1/(k-1)})$, and that it is nonclustered and symmetric. The threshold for van der Waerden's theorem in $[\![N]\!]_p$ is known to lie at $N^{-1/(k-1)}$; this was first proved by Graham, Rödl, and Ruciński in [14], for 3-APs, and later extended by Rödl and Ruciński [25, 26] to general k-APs. The 1-statement implicit in (A3) in our setting is an immediate consequence of this, as every k-AP in $[\![N]\!]$ is also a k-AP in $\mathbb{Z}_N$; independently, it can also be recovered using Proposition 2.11 – robust noncolourability of $\mathcal{H}_{k\text{-AP}}$ follows from van der Waerden's theorem and Varnavides's averaging argument. The 0-statement in (A3) requires extra consideration since $\mathbb{Z}_N$ contains more progressions than $[\![N]\!]$. The arguments of [25, 26] can be echoed here, but we instead prove the stronger Theorem 1.5, which asserts that, when $p < cN^{-1/(k-1)}$ for a sufficiently small constant c, then $(\mathbb{Z}_N)_p$ is a.a.s. list-k-van der Waerden.

This will leave us with verifying the last two assumptions of Theorem 2.1. The choosability assumption (A4) is a simple corollary of Theorem 1.5, see Corollary 7.17 in Section 7.4. Finally, the rainbow star-constellation property follows readily from Szemerédi's theorem and another application of Varnavides's averaging argument. Indeed, suppose that a partial colouring of $\mathbb{Z}_N$ contains $\Omega(N^r)$ rainbow stars. Then, for some $i \in [\![r]\!]$, the set A_i of all elements of $\mathbb{Z}_N$ that are the centres of $\Omega(N^{r-1})$ many

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

i-rainbow stars has $\Omega(N)$ elements. Therefore, A_i contains $\Omega(N^2)$ many k-term APs and each such k-AP is the base of $\Omega(N^{(r-1)k})$ rainbow constellations.

7.3. Schur's theorem

In this section, we prove Theorem 1.4, which asserts that the property of being r-Schur has a sharp threshold in $(\mathbb{Z}_N)_p$. As in the case of van der Waerden's theorem, we have to work in $\mathbb{Z}_N$, as the interval $[\![N]\!]$ lacks the symmetries required by Theorem 2.1. Define the Schur hypergraph $\mathcal{H}$ on $\mathbb{Z}_N$ whose edges are *Schur triples*, that is, triples of distinct $x, y, z \in \mathbb{Z}_N$ such that $x + y = z$. This hypergraph has $\Theta(N^2)$ edges, and thus $p_{\mathcal{H}} = \Theta(N^{-1/2})$, and it is easily seen to be nonclustered. However, it is not symmetric. Indeed, every automorphism of $\mathcal{H}$ is of the form $x \mapsto c \cdot x$ for some invertible $c \in \mathbb{Z}_N$. Since 0 is always mapped to itself, the automorphism group is nontransitive. However, when N is a prime number, this is the only obstruction. In other words, the hypergraph $\mathcal{H}' \coloneqq \mathcal{H} - \{0\}$ is symmetric. This is very fortunate because a.a.s. 0 is not in $(\mathbb{Z}_N)_p$ when $p = o(1)$, and as a result $\mathcal{H}_p$ acts very much like $\mathcal{H}'_p$. Put succinctly, we have the following observation:

Observation 7.15. Given $p = o(1)$ and a property $\mathcal{P}$, the probability of $\mathcal{H}_p \in \mathcal{P}$ tends to 0 if and only if the probability that $\mathcal{H}'_p \in \mathcal{P}$ tends to 0.

We aim to use this observation to prove a sharp threshold for $\mathcal{H}$ by proving it first for the symmetric $\mathcal{H}'$. To do so, we will prove that $\mathcal{H}$, on top of being nonclustered, satisfies assumptions (A3)–(A5). These properties will then transfer to $\mathcal{H}'$ and, together with symmetry, we will be able to argue that it has a sharp threshold.

Regarding (A3), Graham, Rödl, and Ruciński [14] located the threshold for Schur's theorem in $[\![N]\!]_p$ at $N^{-1/2}$. Similarly as in the context of van der Waerden's theorem, since $\mathbb{Z}_N$ has more Schur triples than $[\![N]\!]$, this result implies the 1-statement in (A3), but the 0-statement needs extra work. Instead of adjusting the arguments of [14], we prove the stronger Theorem 1.6, which asserts that $N^{-1/2}$ is a threshold for the 'stronger' property of being list-Schur. This solution has the advantage that the choosability assumption (A4) is a simple corollary of this stronger theorem, see Corollary 7.17 in Section 7.4. This leaves us with establishing the rainbow star-constellation property for the Schur hypergraph, which we do in the remainder of this section.

7.3.1. Rainbow star-constellation property for Schur triples

In this short subsection, we verify that the Schur hypergraph has the rainbow star-constellation property. We write X instead of $\mathbb{Z}_N$, noting that our arguments remain valid if we replace it with an arbitrary Abelian group of order N.

It will be convenient to define structures that offer a slight relaxation of the notions of stars and constellations in that their elements may not be distinct: Given a sequence Y of t subsets of X define a *Y-prestar* to be a pair x, y of sequences of t elements of X together with an element $a \in X$, such that $x_i, y_i \in Y_i$ and x_i, y_i, and z form a sum for every $i \in [\![t]\!]$. The element a is called the *centre* of the prestar and each pair x_i, y_i is called a *ray*. Similarly, define a *Y-preconstellation* to be a triplet of Y-prestars whose centres form a sum. Since fixing any two coordinates of a Y-preconstellation leaves at most $O(N^{3t})$ options for completing it, there are at most $O(N^{3t+1})$ many Y-preconstellations that have a repeating coordinate. In particular, the following statement (with Y being the sequence of some $r-1$ colour classes) implies the rainbow star-constellation property.

Proposition 7.16. *For every $\beta > 0$, there exists a $\gamma > 0$ such that the following holds. For any sequence Y of t subsets of X, if there are βN^{t+1} many Y-prestars, then there are γN^{3t+2} many Y-preconstellations.*

Proof. We may always order the elements of the i^{th} ray of any prestar as x_i, y_i so that one of $x_i \pm y_i$ equals the centre of the prestar. Define the *sign pattern* of the prestar to be the sequence of signs that y_i appeared with for each $i \in [\![t]\!]$. If there are βN^{t+1} many Y-prestars, the pigeonhole principle dictates that there are at least $2^{-t}\beta N^{t+1}$ many Y-prestars sharing a specific sign pattern. Call these prestars the

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

popular Y-prestars. Note that every popular Y-prestar x, y with centre a is uniquely determined by x and a since knowing the sign pattern allows us to compute y.

Let $f(x,a)$ be the indicator of whether $x \in X^t$ and $a \in X$ determine a popular Y-prestar and note that $\sum_{x,a} f(x,a) \geqslant 2^{-t}\beta N^{t+1}$. Next, for an $x \in X^t$ and $a, b \in X$, define $f_1(x,a,b) \coloneqq f(x,a) \cdot f(x,b)$. Using Jensen's inequality we learn that

$$\frac{1}{N^t}\sum_{x,a,b} f_1(x,a,b) = \frac{1}{N^t}\sum_x \left(\sum_a f(x,a)\right)^2 \geqslant \left(\frac{1}{N^t}\sum_{x,a} f(x,a)\right)^2 \geqslant \left(2^{-t}\beta N\right)^2,$$

which implies that $\sum_{x,a,b} f_1(x,a,b) \geqslant 2^{-2t}\beta^2 N^{t+2}$.

Now, for $x, x', x'' \in X^t$ and $a, b \in X$, write

$$f_2(x,x',x'',a,b) \coloneqq f_1(x,a,b) \cdot f_1(x',a,b) \cdot f_1(x'',a,b).$$

Using Jensen's inequality again, we get the bound

$$\begin{aligned}\frac{1}{N^2}\sum_{x,x',x'',a,b} f_2(x,x',x'',a,b) &= \frac{1}{N^2}\sum_{a,b}\left(\sum_x f_1(x,a,b)\right)^3 \\ &\geqslant \left(\frac{1}{N^2}\sum_{a,b,x} f_1(x,a,b)\right)^3 \geqslant \left(2^{-2t}\beta^2 N^t\right)^3,\end{aligned}$$

which implies that $\sum_{x,x',x'',a,b} f_2(x,x',x'',a,b) \geqslant 2^{-6t}\beta^6 N^{3t+2}$.

Now, suppose that $f_2(x,x',x'',a,b) = 1$ for some $x, x', x'' \in X^t$ and $a, b \in X$. Then $f(x',a) = f(x'',b) = f(x,a) = f(x,b) = 1$. By definition, this ensures the existence of the following four popular Y-prestars: x, y and x', y' centred at a and x, z and x'', z'' centred at b (the sequences $y, y', z, z'' \in X^t$ are uniquely determined). In particular, for every $i \in [\![t]\!]$, either $x_i + y_i = a$ and $x_i + z_i = b$ or $x_i - y_i = a$ and $x_i - z_i = b$; this means that either $y_i - z_i = a - b$ or $z_i - y_i = a - b$. Therefore, y, z is a Y-prestar centered at $a - b$ and, consequently, (x', y', a), (x'', z'', b), $(y, z, a - b)$ is a Y-preconstellation.

Note that the function mapping x, x', x'', a, b as above to the Y-preconstellation comprising (x', y', a), (x'', z'', b), and $(y, z, a-b)$ is injective. Indeed, one can reconstruct x given y and a, as we know the sign-pattern of the popular prestar x, y centered at a. As a result, we learn that there are at least $2^{-6t}\beta^6 N^{3t+2}$ many Y-preconstellations. □

7.4. Choosability of almost-linear hypergraphs

In this section, we complete the derivations of Theorems 1.3 and 1.4 by proving Theorem 7.4, which immediately implies Theorems 1.5 and 1.6.

Theorem 1.17. *Suppose that $s \geqslant 3$ and that a sequence of s-uniform hypergraphs $\mathcal{H}$ satisfies $\Delta_2(\mathcal{H}) = O(1)$ and $v(\mathcal{H}) \to \infty$. There is a positive c such that, for every $p \leqslant c \cdot v(\mathcal{H})^{-1/(s-1)}$,*

$$\mathbb{P}\big(\mathcal{H}[V_p] \text{ is 2-choosable}\big) \to 1.$$

Even though this is not the exact statement we need for verifying assumption (A4) in the context of van der Waerden's and Schur's theorem, we will be able to obtain the latter as a straightforward corollary.

Corollary 7.17. *Given $s \geqslant 3$ and a nonclustered s-uniform hypergraph $\mathcal{H}$ with $\Delta_2(\mathcal{H}) = O(1)$, a sequence $p = \Theta(p_{\mathcal{H}})$, and any constant K, the probability that $\mathcal{H}_p$ contains a non-2-choosable set of size K tends to zero.*

Proof. Let $\mathcal{F}$ be the family of all non-2-choosable sets with at most K elements and let $\mu(p) \coloneqq \mathbb{P}(\exists B \in \mathcal{F}\colon B \subseteq V(\mathcal{H})_p)$. Theorem 7.4 tells us that there is some $p = \Theta(p_{\mathcal{H}})$ for which $\mu(p) = o(1)$. By

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

monotonicity, $\mu(p') = o(1)$ also for any $p' < p$. Now, for any constant $C > 1$ we can use Lemma 4.6 to bound $\mu(p) \geqslant C^{-K}\mu(Cp)$, so $\mu(Cp) = o(1)$ as well. □

We will describe a process for revealing connected subsets of the vertices of an s-uniform hypergraph $\mathcal{H}$ by layers. Let $N \coloneqq v(\mathcal{H})$ and identify the vertices of $\mathcal{H}$ with the set $[\![N]\!]$. This labeling induces a total ordering of the vertices and also of the edges of $\mathcal{H}$, via the lexicographic ordering. Given a subset S of the vertices of $\mathcal{H}$ which induces a connected subhypergraph of $\mathcal{H}$ we define the following procedure.

Let v_1 be the smallest vertex in S and write $S' = \{v_1\}$.
for $i = 1, 2, \ldots$ **do**
  `// Degenerate steps`
  while *there exists an edge $e \in \mathcal{H}[S]$ such that $2 \leqslant |e \cap S'| \leqslant s-1$* **do**
    Pick the smallest such edge e and let $S' \leftarrow S' \cup e$.
  `// Finish the layer`
  Write $S_i = S'$.
  if $S_i = S$ **then**
    Terminate.
  else
    `// Start a new layer`
    Take the smallest vertex $v_i \in S_i$ which has an edge $e \in \mathcal{H}[S]$ that intersects S_i exactly in v_i.
    Pick the smallest such edge e_i and let $S' \leftarrow S_i \cup e_i$.

Observe that connectivity of S ensures that the process will indeed terminate. We say that a vertex in S is *degenerate* if it was added in a degenerate step. Let $d(S)$ denote the number d such that $S = S_d$.

We will apply this procedure to sets S that are minimally non-2-choosable. (Note that minimality implies that these sets are connected.) The proof of Theorem 7.4 has two steps. First, using a deterministic argument, we will show that each minimally non-2-choosable set must either have at least $s-1$ degenerate elements or contain a structure which we will call a clot. Second, we will see that, when $\mathcal{H}$ satisfies the assumptions of the theorem, connected sets containing either $s-1$ degenerate elements or a clot are too rare to appear in $V(\mathcal{H})_p$.

Definition 7.18. A set A of $2s-3$ vertices in an s-uniform hypergraph is a *nucleus* if, for every $(s-1)$-element subset $A' \subseteq A$, there are two distinct vertices $v_1(A'), v_2(A') \notin A$ such that $A' \cup \{v_i(A')\}$ is an edge for both $i \in \{1,2\}$. A *clot* around a nucleus A is the union of A together with all the vertices $v_i(A')$.

Lemma 7.19. *If $s \geqslant 3$, then every minimally non-2-choosable set of vertices of an s-uniform hypergraph contains either at least $s-1$ degenerate elements or a clot.*

Proof. Let S be a minimally non-2-choosable set of vertices of an s-uniform hypergraph $\mathcal{H}$. We run the process of revealing S described above and let $d = d(S)$. Since $S_1 = \{v_1\}$ is 2-choosable, d must be greater than 1. We may assume that S contains at most $s-2$ degenerate elements, as otherwise there is nothing left to prove. We will show that $A \coloneqq S \setminus S_{d-1}$ is the nucleus of a clot.

Note that A contains at most $2s-3$ elements: the $s-1$ elements of $e_{d-1} \setminus S_{d-1}$ plus at most $s-2$ additional degenerate elements included in subsequent degenerate steps. Further, note that every edge of S which is not contained in S_{d-1} must contain at least $s-1$ elements from A, since otherwise it would intersect S_{d-1} in at least two elements and would have therefore been absorbed into S_{d-1} in a degenerate step. For that reason, every colouring of S_{d-1} may be extended to S unless we are forced to colour some $s-1$ of A with the same colour. However, this may only happen if $|A| = 2s-3$ and all the colour lists of the elements of A are identical, say they comprise the colours red and blue. Moreover, every set of $s-1$ elements in A must be contained in at least two edges: one with an element of S_{d-1} already coloured red, and a second with an element of S_{d-1} already coloured blue. Therefore, A must be the nucleus of a clot. □

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Lemma 7.20. *Suppose that $s \geqslant 3$ and $\mathcal{H}$ is a sequence of s-uniform hypergraphs with $\Delta_2(\mathcal{H}) = O(1)$. If $p = O\big(v(\mathcal{H})^{-1/(s-1)}\big)$, then a.a.s. $\mathcal{H}_p$ does not contain a clot.*

Proof. As before, we write $N := v(\mathcal{H})$ and identify $V(\mathcal{H})$ with $[\![N]\!]$. We give two different arguments, depending on whether or not $s > 3$.

We first show that, when $s > 3$, there are only $O(N^2)$ clots. This will be sufficient as every clot contains at least $2s - 1$ vertices (the nucleus and at least two additional vertices) and thus the expected number of clots in $\mathcal{H}_p$ is $O(N^2 p^{2s-1}) = O(p) = o(1)$. There are at most $O(N^2)$ ways to choose the two smallest elements of the nucleus A of a clot. Since $s - 1 > 2$, every other element of A belongs to an edge containing these two elements and there are only $\Delta_2(\mathcal{H}) = O(1)$ such edges. Similarly, every element of the clot that is not in A belongs to one of the at most $\binom{2s-3}{s-1} \cdot \Delta_{s-1}(\mathcal{H}) = O(1)$ edges that intersect A in $s - 1$ elements.

The second case is where $s = 3$, implying that every nucleus has three vertices and every clot has at least five vertices. In particular, there are at most N^3 ways to choose the nucleus A of a clot and every element of the clot that is not in A belongs to one of the at most $3 \cdot \Delta_2(\mathcal{H}) = O(1)$ edges that intersect A in two vertices. Consequently, the expected number of clots with at least seven vertices is $O(N^3 p^7) = O(p) = o(1)$. If a clot has fewer than seven vertices, then, by the pigeonhole principle, there must be some vertex v not in its nucleus that forms edges with two different pairs of vertices from the nucleus. This implies that there are only $O(N^2)$ such clots: We may pick one such v and an element of the nucleus with at most N^2 options. The remaining vertices of the clot can be added one-by-one in such a way that the added element forms an edge with two previously added vertices. This means that the expected number of such clots is $O(N^2 p^5) = O(p) = o(1)$. □

Lemma 7.21. *Suppose that $s \geqslant 3$ and $\mathcal{H}$ is a sequence of s-uniform hypergraphs with $\Delta_2(\mathcal{H}) = O(1)$. If $p \leqslant cv(\mathcal{H})^{-1/(s-1)}$ for a sufficiently small positive constant c, then a.a.s. any connected $S \subseteq V(\mathcal{H})_p$ contains at most $s - 2$ degenerate elements.*

Proof. As before, we write $N := v(\mathcal{H})$ and identify $V(\mathcal{H})$ with $[\![N]\!]$. If a connected set S has at least $s - 1$ degenerate elements, then it must contain a connected subset with at least $s - 1$ degenerate elements for which our procedure executed at most $s - 1$ degenerate steps. Indeed, one obtains such subset by simply halting the procedure after $s - 1$ degenerate elements are revealed, while noting that every degenerate step must introduce at least one new degenerate element. We may thus restrict our attention to connected sets S with this additional property. We further claim that $d(S) \leqslant \frac{|S|-1}{s-1}$ for every S with at least $s - 1$ degenerate elements. To see this, note that the number of degenerate elements in S is $|S| - 1 - (d(S) - 1)(s - 1)$, as precisely $(d(S) - 1)(s - 1)$ vertices are added in nondegenerate steps. Summarising, it is enough to show that the expected number of connected sets S with $d(S) \leqslant \frac{|S|-1}{s-1}$ and at most $s - 1$ degenerate steps that appear in $V(\mathcal{H})_p$ tends to zero.

Let X_k be the number of such sets that have exactly k elements. We can bound X_k using the following logic: First, we fix integers $d \leqslant \frac{k-1}{s-1}$ and $d' \leqslant s - 1$ and bound the number of k-element sets S with $d(S) = d$ for which the procedure runs d' degenerate steps. Given d and d', we decide, for each of the $d + d' - 1$ steps of the procedure revealing S whether it is degenerate or not; there are at most $2^{d+d'-1} \leqslant 2^k$ options. Second, we pick $v_1 \in [\![N]\!]$. Third, for every degenerate step, we choose some two vertices of S' that witnessed $|e \cap S'| \geqslant 2$ and choose the edge e; there are at most $k^2 \cdot \Delta_2(\mathcal{H})$ options. Fourth, we choose the number of times each element of S plays the role of v_i to start a new layer. We may represent this as a multiset M of $[\![k]\!]$ with $d - 1$ elements, where $j \in [\![k]\!]$ corresponds to the jth vertex in the order of arrival to S; thus, there are at most $\binom{k+d-1-1}{d-1} \leqslant 2^{2k}$ options. Crucially, note that, for each $i \in [\![d - 1]\!]$, the identity of the vertex v_i is determined by M and S_i. Finally, we pick, for every $i \in [\![d - 1]\!]$, the edge e_i that intersects S_i in v_i; there are at most $N \cdot \Delta_2(\mathcal{H})$ options. Summarising,

$$X_k \leqslant \sum_{d \leqslant \frac{k-1}{s-1}} \sum_{d' \leqslant s-1} 2^k \cdot N \cdot \big(k^2 \cdot \Delta_2(\mathcal{H})\big)^{d'} \cdot 2^{2k} \cdot \big(N \cdot \Delta_2(\mathcal{H})\big)^{d-1} \leqslant 2^{Ck} \cdot N^{\frac{k-1}{s-1}},$$

where C is a constant that depends only on s and the constant implicit in the upper bound $\Delta_2(\mathcal{H}) = O(1)$.

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Suppose now that $p \leqslant 2^{-C}N^{-1/(s-1)}$ and let $\mathcal{B}$ denote the event that $[\![N]\!]_p$ contains a connected subset with at least $s-1$ degenerate elements. We have

$$\mathbb{P}(\mathcal{B}) \leqslant \sum_{k\geqslant 1} X_k \cdot p^k \leqslant \sum_{k\geqslant 1} 2^{Ck} \cdot N^{\frac{k-1}{s-1}} \cdot 2^{-Ck} N^{-\frac{k}{s-1}} = N^{-\frac{1}{s-1}} = o(1),$$

as claimed. □

A. Weak threshold – the 1-statement

We begin by stating two weaker versions of the notion of nonclustered hypergraphs.

Definition A.1. A sequence of nonempty, s-uniform hypergraphs $\mathcal{H}$ is *weakly nonclustered* if

$$\Delta_t(\mathcal{H}) = O\left(p_{\mathcal{H}}^{t-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}\right) \qquad \text{for all } t \in [\![s-1]\!].$$

We say that it is *mildly nonclustered* if it is weakly nonclustered and, additionally, $\Delta_u(\mathcal{H}) \ll p_{\mathcal{H}}^{u-1} \cdot e(\mathcal{H})/v(\mathcal{H})$ for some $u \in \{2,\dots,s-1\}$.

Proposition A.2. *Suppose that $\mathcal{H}$ is weakly nonclustered and robustly non-r-colourable. Then there are constants C, α such that, for any $p > C \cdot p_{\mathcal{H}}$, we have*

$$\mathbb{P}\big(\mathcal{H}_p \text{ is } r\text{-colourable}\big) \leqslant \exp(-\alpha \cdot p \cdot v(\mathcal{H})).$$

In particular, if $v(\mathcal{H}) \cdot p_{\mathcal{H}}$ tends to infinity, then the proposition tells us that the probability of $\mathcal{H}_p$ being r-colourable tends to zero. Before turning to the proof, we pause to make a couple of remarks concerning this point.

Remark. First, we claim that when $\mathcal{H}$ satisfies the stronger assumption of being mildly nonclustered, then $v(\mathcal{H}) \cdot p_{\mathcal{H}}$ does tend to infinity. Indeed, if there is an index $u \in \{2,\dots,s-1\}$ such that $\Delta_u(\mathcal{H}) \ll p_{\mathcal{H}}^{u-1} \cdot e(\mathcal{H})/v(\mathcal{H})$, then, bounding $\Delta_u(\mathcal{H})$ from below by the average degree of a u-element subset of $V(\mathcal{H})$, we can write

$$\frac{e(\mathcal{H})}{v(\mathcal{H})^u} \leqslant \Delta_u(\mathcal{H}) \ll p_{\mathcal{H}}^{u-1} v(\mathcal{H})^{u-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})^u}$$

and conclude that $p_{\mathcal{H}} \cdot v(\mathcal{H}) \gg 1$, as requested.

Remark. Next, it is possible for a sequence of hypergraphs $\mathcal{H}$ and probabilities p fulfilling the requirements of Proposition A.2 that the probability that $\mathcal{H}_p$ is r-colourable to be bounded away from zero. Indeed, let $\mathcal{H}$ be the complete s-uniform hypergraph on n vertices. Since $e(\mathcal{H}) = \Theta(n^s)$, we have that $p_{\mathcal{H}} = \Theta(n^{-1})$, and therefore that $v(\mathcal{H}) \cdot p_{\mathcal{H}} = \Theta(1)$.

We verify that $\mathcal{H}$ satisfies the assumptions of the proposition. First, for every $t \in [\![s]\!]$, we have $\Delta_t(\mathcal{H}) = \Theta(n^{s-t}) = \Theta\big(p^{t-1} \cdot e(\mathcal{H})/v(\mathcal{H})\big)$, so $\mathcal{H}$ is weakly nonclustered. Second, by the pigeonhole principle, every r-colouring of $V(\mathcal{H})$ must have at least n/r elements sharing the same colour. These vertices alone induce $\Theta(n^s)$ monochromatic edges.

However, for every $C > 0$ and $p = C \cdot p_{\mathcal{H}}$, the probability that $\mathcal{H}_p$ is in fact empty is $(1-p)^{v(\mathcal{H})} \geqslant \exp(-\Theta(v(\mathcal{H}) \cdot p)) = \exp(-\Theta(1))$. This of course implies that the probability that $\mathcal{H}_p$ is r-colourable is bounded away from zero.

Proof of Proposition A.2. Let δ be the constant for which $\mathcal{H}$ is robustly non-r-colourable. Since $\mathcal{H}$ is weakly nonclustered, we may apply to it the Container Lemma (Theorem B). Write $V = V(\mathcal{H})$. Taking any $\varepsilon < \delta/r$, there are constants $c_1 = c_1(\varepsilon)$, $t = t(\varepsilon)$, and a function $f\colon \mathcal{P}(V)^t \to \mathcal{P}(V)$ such that the following hold:

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

(1) For every independent set I there are $T_1, \dots, T_t \subseteq I$ with at most $c_1 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})$ elements, such that $I \subseteq f(T_1, \dots, T_t)$.
(2) The set $f(T_1, \dots, T_t)$ induces at most $\varepsilon e(\mathcal{H})$ edges in $\mathcal{H}$.

Suppose now that $\mathcal{H}_p$ is r-colourable. This of course means that V_p is a union of r independent sets $I_1, \dots, I_r$. Following the Container Lemma, this implies that, for each $i \in [\![r]\!]$, there are sets $T_1^i, \dots, T_t^i \subseteq I_i$, each containing at most $c_1 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})$ vertices, such that $I_i \subseteq C_i := f(T_1^i, \dots, T_t^i)$. Of course, that would mean that all subsets T_j^i are contained in $V(\mathcal{H}_p)$ and that $V(\mathcal{H}_p)$ is covered by the containers $C_1, \dots, C_r$. Using the fact that the containers induce few edges, together with the robust noncolourability, we will be able to prove the following claim.

Claim A.3. *There is a constant α such that the number of vertices in $V \setminus (C_1 \cup \cdots \cup C_r)$ is at least $\alpha v(\mathcal{H})$.*

Proof of the claim. Write $W := V \setminus (C_1 \cup \cdots \cup C_r)$ and consider the following r-colouring of the vertices of $\mathcal{H}$. Colour each vertex according to the index i of its container C_i (if it is contained in more then one container, choose one arbitrarily), and colour the vertices of W by the colour 1. Since $\mathcal{H}$ is robustly non-r-colourable, we know that there are at least $\delta e(\mathcal{H})$ monochromatic edges. Since every container induces at most $\varepsilon e(\mathcal{H})$ edges, there must be at least $(\delta - r\varepsilon) \cdot e(\mathcal{H})$ monochromatic edges in the colour 1 that have at least one vertex in W. On the other hand, there are at most $\Delta_1(\mathcal{H}) \cdot |W|$ such edges. By our assumption on $\mathcal{H}$, there is a constant K such that

$$(\delta - r\varepsilon) \cdot e(\mathcal{H}) \leqslant \Delta_1(\mathcal{H}) \cdot |W| \leqslant K \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \cdot |W|.$$

Remembering that $\varepsilon < \delta/r$, there must be a constant $\alpha > 0$ such that $|W| > \alpha v(\mathcal{H})$, as promised. □

Following the previous discussions, and letting $c_2 = rtc_1$, we can bound

$$\begin{aligned}
&\mathbb{P}\big(\mathcal{H}_p \text{ is } r\text{-colourable}\big) \\
&\leqslant \sum_{|T_j^i| \leqslant c_1 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})} \mathbb{P}\left(\bigcup_{j \in [\![t]\!], i \in [\![r]\!]} T_j^i \subseteq V_p \subseteq C_1 \cup \cdots \cup C_r \right) \\
&= \sum_{|T_j^i| \leqslant c_1 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})} \mathbb{P}\left(\bigcup T_j^i \subseteq V_p \right) \cdot \mathbb{P}\left(V_p \cap (V \setminus (C_1 \cup \cdots \cup C_r)) = \emptyset \right) \\
&\leqslant \sum_{k \leqslant rtc_1 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})} \binom{v(\mathcal{H})}{k} \cdot 2^{rtk} \cdot p^k \cdot (1-p)^{\alpha v(\mathcal{H})} \\
&\leqslant c_2 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H}) \cdot \max_{k \leqslant c_2 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})} \left(\frac{2^{rt} e v(\mathcal{H}) p}{k} \right)^k \cdot \exp(-\alpha v(\mathcal{H}) \cdot p).
\end{aligned}$$

The function $k \mapsto \left(\frac{eM}{s}\right)^k$ is increasing until it reaches its maximum at $k = M$. Therefore, supposing that C is large enough so that $c_2 < 2^{rt} \cdot C$, the maximum is achieved at $s = c_2 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})$, allowing us to bound

$$\begin{aligned}
\max_{k \leqslant c_2 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H})} \left(\frac{2^{rt} e \cdot v(\mathcal{H}) \cdot p}{k} \right)^k &\leqslant \left(\frac{2^{rt} e \cdot v(\mathcal{H}) \cdot p}{c_2 \cdot v(\mathcal{H}) \cdot p_{\mathcal{H}}} \right)^{c_2 \cdot v(\mathcal{H}) \cdot p_{\mathcal{H}}} \\
&= \exp\left(c_2 \log\left(\frac{2^{rt} e}{c_2} \cdot \frac{p}{p_{\mathcal{H}}} \right) \cdot \frac{p_{\mathcal{H}}}{p} \cdot v(\mathcal{H}) \cdot p \right).
\end{aligned}$$

Recalling that $C \leqslant \frac{p}{p_{\mathcal{H}}}$ and $x^{-1} \log x \to 0$ as x tends to infinity, taking C sufficiently large, we have $c_2 \log\left(\frac{2^{rt} e}{c_2} \cdot \frac{p}{p_{\mathcal{H}}}\right) \cdot \frac{p_{\mathcal{H}}}{p} < \frac{1}{2}\alpha$ and, therefore,

$$\mathbb{P}\big(\mathcal{H}_p \text{ is } r\text{-colourable}\big) \leqslant c_2 \cdot p_{\mathcal{H}} \cdot v(\mathcal{H}) \cdot \exp\left(-\frac{1}{2}\alpha \cdot p \cdot v(\mathcal{H}) \right) \leqslant \exp(-\alpha' \cdot p \cdot v(\mathcal{H})).$$ □

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

B. A container lemma for sparse sets

Theorem 4.7. *For every positive integer k and all positive reals ε and K, there exist an integer t and a positive real δ such that the following holds. Suppose that a nonempty k-uniform (multi)hypergraph $\mathcal{G}$ with vertex set V and a positive real τ satisfy*

$$\Delta_\ell(\mathcal{G}) \leqslant K\tau^{\ell-1} \cdot \frac{e(\mathcal{G})}{v(\mathcal{G})}$$

for every $\ell \in [\![k]\!]$. Then, there exists a function $f\colon \mathcal{P}(V)^t \to \mathcal{P}(V)$ with the following properties:

(i) *For every set $I \subseteq V$ satisfying $e(\mathcal{G}[I]) \leqslant \delta\tau^k e(\mathcal{G})$, there are $S_1, \dots, S_t \subseteq I$ with at most $\tau v(\mathcal{G})$ elements each such that $I \subseteq f(S_1, \dots, S_t)$.*
(ii) *For every $S_1, \dots, S_t \subseteq V$, the set $f(S_1, \dots, S_t)$ induces fewer than $\varepsilon e(\mathcal{G})$ edges in $\mathcal{G}$.*

Theorem B.1 (27, Corollary 3.6). *For every positive integer k and positive real ε, there exists an integer s such that the following holds. Suppose that a nonempty k-uniform (multi)hypergraph $\mathcal{G}$ and $\tau \in (0, 1/2)$ satisfy*

$$\delta(\mathcal{G}, \tau) := 2^{\binom{k}{2}-1} \sum_{j=2}^{k} 2^{-\binom{j-1}{2}} \delta_j(\mathcal{G}, \tau) \leqslant \frac{\varepsilon}{12k!},$$

where

$$\delta_j(\mathcal{G}, \tau) := \frac{\tau^{1-j}}{ke(\mathcal{G})} \cdot \sum_{v \in V} \max\{\deg T : v \in T \subseteq V \text{ and } |T| = j\}.$$

Then there exists a function $C\colon \mathcal{P}(V)^s \to \mathcal{P}(V)$ such that, letting

$$\mathcal{T} := \big\{(T_1, \dots, T_s) \in \mathcal{P}(V)^s : |T_i| \leqslant s\tau|V| \text{ for all } i \in [\![s]\!]\big\},$$

we have:

(a) *For every set $I \subseteq V$ satisfying $e(\mathcal{G}[I]) \leqslant 24\varepsilon k! k\tau^k e(\mathcal{G})$, there exists $T = (T_1, \dots, T_s) \in \mathcal{T} \cap \mathcal{P}(I)^s$ with $I \subseteq C(T)$.*
(b) *For every $T \in \mathcal{T}$, the set $C(T)$ induces at most $\varepsilon e(\mathcal{G})$ edges in $\mathcal{G}$.*

Derivation of Theorem B from Theorem B.1. Let $\mathcal{G}$ be a nonempty k-uniform hypergraph with vertex set V and let ε and K be positive reals. We set $s := s_{B.1}(k, \varepsilon)$ and let

$$L := \left\lceil \frac{12(k-1)!2^{\binom{k}{2}}K}{\varepsilon} \right\rceil, \qquad t := 2L^2 s, \qquad \text{and} \qquad \delta := 24\varepsilon k! k L^k$$

Suppose that the maximum degrees of $\mathcal{G}$ satisfy the assumptions of the theorem for some τ. Note that, for every $j \in \{2, \dots, k\}$,

$$\delta_j(\mathcal{G}, L\tau) = \frac{(t\tau)^{1-j}}{ke(\mathcal{G})} \cdot v(\mathcal{G})\Delta_j(\mathcal{G}) \leqslant \frac{KL^{1-j}}{k}$$

and thus, as $L \geqslant 2$,

$$\delta(\mathcal{G}, L\tau) \leqslant 2^{\binom{k}{2}-1} \cdot \sum_{j=2}^{k} \frac{Kt^{1-j}}{k} \leqslant \frac{2^{\binom{k}{2}}K}{kL} \leqslant \frac{\varepsilon}{12k!}.$$

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

Consequently, Theorem B.1, invoked with $\tau_{B.1} = L\tau$, implies that there exist a function $C\colon \mathcal{P}(V)^s \to \mathcal{P}(V)$ that satisfies (a) and (b). We define $f\colon \mathcal{P}(V)^t \to \mathcal{P}(V)$ by letting

$$f(S_1, \dots, S_t) \coloneqq C\big(S_1 \cup \dots \cup S_{t/s}, \dots, S_{(s-1)t/s+1} \cup \dots \cup S_t\big).$$

In particular, if $I \subseteq V$ satisfies

$$e(\mathcal{G}[I]) \leqslant \delta\tau^k e(\mathcal{G}) \leqslant 24\varepsilon k!k(L\tau)^k e(\mathcal{G}),$$

then there are $T_1, \dots, T_s \subseteq I$, with $|T_i| \leqslant sL\tau|V|$ for each i, such that $I \subseteq C(T_1, \dots, T_s)$. This gives the assertion of the theorem, as we may partition each T_i into t/s sets $S_{t(i-1)/s+1}, \dots, S_{t(i+1)/s}$, each of size at most $\lceil s/t \cdot sL\tau|V| \rceil \leqslant \tau|V|$. □

C. A graph without the star-constellation property

In this appendix we provide an example of a graph whose corresponding hypergraph does not have the rainbow star-constellation property. We need not look further than the usual suspect, the Petersen graph.

Claim C.1. *The hypergraph corresponding to the Petersen graph does not have the rainbow star-constellation property.*

Proof. It is enough to show this for $r = 2$. Note that since the Petersen graph is edge-transitive there is only one star S up to isomorphism and therefore only one constellation C. To show that there is no homomorphism $C \to S$ we will first claim that every homomorphism $S \to S$ must be an isomorphism. To see why this helps us, observe that every homomorphism $S \to S$ must then map the centre vertices to themselves, as they are the only vertices of degree 2. Therefore, any homomorphism $C \to S$, which induces a homomorphism from each of its stars to S, must send all the centre vertices of C to the two centre vertices of S. Viewing this as a 2-colouring of the centre vertices of C, we may use the fact that the Petersen graph is not 2-colourable to find a star in C where both of the centre vertices were mapped to the same vertex. However, this would mean that the homomorphism from that star was not an isomorphism, in contradiction to the previous claim.

Suppose $\varphi\colon S \to S$ is a homomorphism. To prove that it is also an isomorphism we show it is injective. Since the Petersen graph has no cycles of length 3, we learn that every 5-cycle must be mapped to a 5-cycle. As a corollary, we learn that whenever $x, y \in S$ are in a 5-cycle, $\varphi(x) \neq \varphi(y)$. So we only need to worry about pairs of elements that are not in a 5-cycle. However, one can verify that there are only five pairs of vertices not in a 5-cycle. Using the labels from Figure 2, these are the centre vertices $\{u, u'\}$, the pairs $\{u, a\}$ and $\{u, b\}$, and by symmetry the pairs $\{u', a'\}$ and $\{u', b'\}$. Of course if φ were to map two vertices x, y to the same vertex z, then it would have to send all neighbours of x and y to the neighbourhood of z. In all the above cases the pair x, y has at least 4 neighbours of which

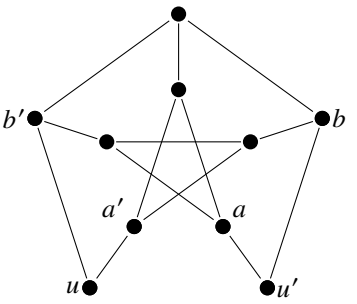


Figure 2. *A Petersen star for $r = 2$.*

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press

every pair is in a 5-cycle, so they cannot be mapped to the same vertex, meaning that z would have to be of degree $\geqslant 4$; however, all vertices in S have degree at most 3. □

Acknowledgements. We thank the anonymous referees for their careful reading and helpful feedback.

References

[1] N. Alon, M. Bucić, T. Kalvari, E. Kuperwasser and T. Szabó, 'List Ramsey numbers', *J. Graph Theory* **96**(1) (2021), 109–128.
[2] J. Balogh, R. Morris and W. Samotij, 'Independent sets in hypergraphs', *J. Amer. Math. Soc.* **28**(3) (2015), 669–709.
[3] J. Balogh and W. Samotij, 'An efficient container lemma', *Discrete Anal.* (2020), Paper No. 17, 56.
[4] B. Bollobás and A. Thomason, 'Threshold functions', *Combinatorica* **7**(1) (1987), 35–38.
[5] P. Erdős and A. Rényi, 'On the evolution of random graphs', *Magyar Tud. Akad. Mat. Kutató Int. Közl.* **5** (1960), 17–61.
[6] J. Fox, X. He, S. Luo and M. Wenqiang Xu, 'Multicolor list Ramsey numbers grow exponentially', *J. Graph Theory* **101**(3) (2022), 389–396.
[7] P. Frankl and V. Rödl, 'Large triangle-free subgraphs in graphs without K_4', *Graphs Combin.* **2**(2) (1986), 135–144.
[8] E. Friedgut, 'Sharp thresholds of graph properties, and the k-sat problem', *J. Amer. Math. Soc.* **12**(4) (1999), 1017–1054, With an appendix by Jean Bourgain.
[9] E. Friedgut, 'Hunting for sharp thresholds', *Random Structures Algorithms* **26**(1–2) (2005), 37–51.
[10] E. Friedgut, H. Hàn, Y. Person and M. Schacht, 'A sharp threshold for van der Waerden's theorem in random subsets', *Discrete Anal.* (2016), Paper No. 7, 20.
[11] E. Friedgut and M. Krivelevich, 'Sharp thresholds for certain Ramsey properties of random graphs', *Random Structures Algorithms* **17**(1) (2000), 1–19.
[12] E. Friedgut, V. Rödl, A. Ruciński and P. Tetali, 'A sharp threshold for random graphs with a monochromatic triangle in every edge coloring', *Mem. Amer. Math. Soc.* **179**(845) (2006), vi+66.
[13] E. Friedgut, V. Rödl and M. Schacht, 'Ramsey properties of random discrete structures', *Random Structures Algorithms* **37**(4) (2010), 407–436.
[14] R. Graham, V. Rödl and A. Ruciński, 'On Schur properties of random subsets of integers', *J. Number Theory* **61**(2) (1996), 388–408.
[15] T. Harris, 'A lower bound for the critical probability in a certain percolation process', *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 56 (Cambridge University Press, 1960), 13–20.
[16] S. Janson, 'Poisson approximation for large deviations', *Random Structures Algorithms* **1**(2) (1990), 221–229.
[17] E. Kuperwasser and W. Samotij, 'The list-Ramsey threshold for families of graphs', *Combinatorics, Probability and Computing* **33**(6) (2024), 829–851.
[18] E. Kuperwasser, W. Samotij and Y. Wigderson, 'On the Kohayakawa–Kreuter conjecture', *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 178 (Cambridge University Press, 2025), 293–320.
[19] T. Łuczak, A. Ruciński and B. Voigt, 'Ramsey properties of random graphs', *J. Combin. Theory Ser. B* **56**(1) (1992), 55–68.
[20] R. Nenadov, Y. Person, N. Škorić and A. Steger, 'An algorithmic framework for obtaining lower bounds for random Ramsey problems', *J. Combin. Theory Ser. B* **124** (2017), 1–38.
[21] R. Nenadov and A. Steger, 'A short proof of the random Ramsey theorem', *Combin. Probab. Comput.* **25**(1) (2016), 130–144.
[22] F.P. Ramsey, 'On a problem of formal logic', *Proceedings of the London Mathematical Society* **2**(1) (1930), 264–286.
[23] V. Rödl and A. Ruciński, *Lower bounds on probability thresholds for Ramsey properties*, Combinatorics, Paul Erdős is eighty, Vol. 1, Bolyai Soc. Math. Stud., János Bolyai Math. Soc., Budapest, 1993, pp. 317–346.
[24] V. Rödl and A. Ruciński, 'Random graphs with monochromatic triangles in every edge coloring', *Random Structures Algorithms* **5**(2) (1994), 253–270.
[25] V. Rödl and A. Ruciński, 'Threshold functions for Ramsey properties', *J. Amer. Math. Soc.* **8**(4) (1995), 917–942.
[26] V. Rödl and A. Ruciński, 'Rado partition theorem for random subsets of integers', *Proc. London Math. Soc. (3)* **74**(3) (1997), 481–502.
[27] D. Saxton and A. Thomason, 'Hypergraph containers', *Invent. Math.* **201**(3) (2015), 925–992.
[28] M. Schacht and F. Schulenburg, 'Sharp thresholds for Ramsey properties of strictly balanced nearly bipartite graphs', *Random Structures Algorithms* **52**(1) (2018), 3–40.
[29] F. Schulenburg, *Threshold results for cycles*, Ph.D. thesis, University of Hamburg, 2016.
[30] I. Schur, 'Über die Kongruenz $x^m + y^m \equiv z^m \pmod p$', *Jahresbericht der Deutschen Mathematiker-Vereinigung* **25** (1917), 114–116.
[31] B.L. van der Waerden, 'Beweis einer Baudetschen Vermutung', *Nieuw Arch. Wiskd., II. Ser.* 15 (1927), 212–216 (German).

https://doi.org/10.1017/fms.2026.10178 Published online by Cambridge University Press